arXiv:1703.03412v1 [math.ST] 9 Mar 2017

Uniform estimation of a class of random graph functionals

Ismaël Castillo and Peter Orbanz

Université Pierre et Marie Curie – Paris VI
Laboratoire Probabilités et Modèles Aléatoires
4, place Jussieu
75005 Paris, France
&
Columbia University
Department of Statistics
1255 Amsterdam Avenue
New York, NY 10027, USA
e-mail: ismael.castillo@upmc.fr; porbanz@stat.columbia.edu

Abstract: We consider estimation of certain functionals of random graphs. The random graph is generated by a stochastic block model (SBM). The number of classes is fixed or grows with the number of vertices. Minimax lower and upper bounds of estimation along specific submodels are derived. The results are nonasymptotic and imply that uniform estimation of a single connectivity parameter is much slower than the expected asymptotic pointwise rate. Specifically, the uniform quadratic rate does not scale as the number of edges, but only as the number of vertices. The lower bounds are local around any possible SBM. An analogous result is derived for functionals of a class of smooth graphons.

MSC 2010 subject classifications: Primary 62G20.
Keywords and phrases: Stochastic Blockmodel, semiparametric estimation of functionals, minimax rates, graphon model.

1. Introduction

Network data occurs in a range of fields, and its analysis has become a highly interdisciplinary effort [14, 18, 24, 26]. In statistical network analysis, two classes of models have recently received particular attention: Graphon models [9, 20, 27], and the subclass of stochastic block models (SBMs) [1, 7, 17]. These models parametrize a random graph by a symmetric measurable function w, which can be interpreted as representing an adjacency matrix in the limit of infinite graph size [9]. In a SBM, the function is in particular piece-wise constant. Examples of statistical problems arising in this field include *estimation problems* (see below), *class label recovery* [e.g. 5, 28, 32], and *signal detection*, which refers to testing for the presence of a signal in settings where observed data constitutes a network or array [e.g. 3, 4, 10, 29].

The work described here concerns estimation problems. SBMs are parametrized by the sizes of communities, which subdivide the domain of w into regions of constant value, and by the values of the function w on these regions. The estimation of these parameters is typically complicated by the fact that the specific subdivision in regions, or 'labelling', is not observed. Remarkably, the pointwise asymptotic results obtained in [7] imply SBM parameters can be recovered at the same rate regardless of whether or not labels are observed (at least asymptotically and on the interior of the parameter set). Earlier results include [1, 11, 12]. Other functionals can also be of interest; estimation of moments, for instance, is considered in [2, 8]. In more general graphon models, one can estimate the complete function w [16, 23, 30]. Of particular relevance for our purposes is an idea highlighted by Gao, Lu, and Zhou [16]: If one estimates the entire parameter function w, not observing labels slows the rate—which is not the case in the setting studied in [7].

In this paper, we consider uniform, non-asymptotic rates of convergence for estimation of certain parameters of graphon models. We begin with one of the simplest graphon models conceivable, a two-class SBM specified by a single, scalar parameter. Even under this simple model, uniform and pointwise rates differ significantly. We then derive results for SBMs with k classes, where k is not

necessarily fixed and may grow to infinity with n. The phenomenon observed for $k = 2$ persists in the general case, with a rate depending on k. Finally, we derive results for certain functionals under a smooth graphon model. Before we describe our results more precisely, we define the models and introduce the notation used in the sequel.

Stochastic blockmodels and graphon models. Consider sampling at random an undirected, simple graph G on the vertex set $\mathcal{V} = \{1, \dots, n\}$ as follows. Fix some $k \in \{1, \dots, n\}$. Let $\pi = (\pi_1, \dots, \pi_k)$ be a probability distribution on the set $\{1, \dots, k\}$, with π identified as a line vector of size k. Let $M := (M_{lm})$ be a symmetric $k \times k$ matrix with elements $M_{lm} \in [0, 1]$. To sample a graph G, we generate its adjacency matrix $X = (X_{ij})_{i,j \in \mathcal{V}}$. Since G is undirected, it suffices to sample entries with $i < j$, and writing $\mathrm{Be}(p)$ as a shorthand for $\mathrm{Bernoulli}(p)$,

1. For each vertex $i \in \mathcal{V}$, independently generate a label $\varphi(i) \sim \pi$.
2. For each pair $i < j$ in $\mathcal{V}$, independently sample $X_{ij} \,|\, \varphi(i), \varphi(j) \sim \mathrm{Be}(M_{\varphi(i)\varphi(j)})$.

In this notation, φ is a (random) mapping $\varphi : \{1, \dots, n\} \to \{1, \dots, k\}$ that attributes a label to each node of the graph. It is random because labels are by definition randomly sampled. The distribution $P_{\pi,M}$ so defined on the set of undirected, simple graphs is called a *stochastic blockmodel* of order k with parameters π and M. One can also write

$$\begin{aligned} (\varphi(1), \dots, \varphi(n)) &\sim \pi^{\otimes n} \\ (X_{ij})_{i<j} \mid \varphi &\sim \bigotimes_{i<j} \mathrm{Be}(M_{\varphi(i)\varphi(j)}), \end{aligned} \tag{1}$$

where $\otimes$ denotes a tensor product of distributions, $\pi^{\otimes n} = \pi \otimes \cdots \otimes \pi$, and here and in the sequel $i < j$ refers to all pairs of indices $(i, j) \in \mathcal{V}^2$ with $i < j$. Any given φ partitions the vertex set $\{1, \dots, n\}$ into k distinct classes, often called *communities*. We call π the *proportions* vector and M the matrix of *connectivity* parameters.

These models can be regarded as a special case of a more general class of random graphs, parametrized by the set of all measurable functions $w : [0, 1]^2 \to [0, 1]$ that are symmetric, i.e. $w(x, y) = w(y, x)$. Any such w defines a random graph G: denoting by $\mathrm{Unif}[0, 1]$ the uniform distribution on $[0, 1]$, and $(U_i)_i = (U_i)_{1 \le i \le n}$, set

$$\begin{aligned} (U_i)_i &\sim \mathrm{Unif}[0, 1]^{\otimes n} \\ (X_{ij})_{i<j} \mid (U_i)_i &\sim \bigotimes_{i<j} \mathrm{Be}(w(U_i, U_j)). \end{aligned} \tag{2}$$

The law P_w of the graph G defined by the random matrix X in (2) is called a *graphon model* [9].

Mixture interpretation. The n-tuple (U_i) in a graphon model, or, equivalenly, the mapping φ in a SBM, are in general not observed, and can hence be interpreted as latent variables. In other words, the distribution of the data $(X_{ij})_{i<j}$ is a *mixture*. Also, either class of models generalises to directed graphs, by dropping the symmetry constraints on π and w, and requiring only $i \neq j$ rather than $i < j$; in the following, we consider only the undirected case.

The data distribution in (2) remains invariant if w is replaced by $w \circ g$, for any measure-preserving transformation g of $[0, 1]^2$: $P_w = P_{w \circ g}$. More generally, two graphons w and w' are considered equivalent if $P_w = P_{w'}$. The equivalence class $\langle w \rangle$ of w is called a graph limit. Similarly in (1), if σ is a fixed arbitrary permutation of $\{1, \dots, k\}$, with permutation matrix Σ, then $P_{\pi,M} = P_{\pi\Sigma, \Sigma M \Sigma^T}$. The parameters of the SBM can only be recovered up to label switching. We refer to [1] and [11] for detailed identifiability statements.

Stochastic blockmodels are special cases of graphon models, obtained by choosing w as a histogram: subdivide the unit interval into k intervals $I_s := [\sum_{i<s} \pi_i, \sum_{i \le s} \pi_i)$ of respective lengths π_s, and set

$$w(x, y) := M_{ij} \quad \text{for} \quad x \in I_i, y \in I_j \,. \tag{3}$$

Then $P_w = P_{\pi,M}$. In a graphon model, the continuous vertex labels U_i are almost surely distinct; in a stochastic block model, labels coincide whenever two vertices belong to the same class. Thus,

the SBM labels can be regarded as discretization of graphon labels. Conversely, any graphon can be approximated by a sequence of stochastic blockmodels of increasing order k; indeed, the set of stochastic blockmodels—that is, of graphons of the form (3) for all k, π and M—is dense in the set of w's endowed with its natural topology [21]. This idea can be used to construct SBM-valued estimators for graphons [30], [16].

Fixed and random design. In models (1)-(2), the latent variables, respectively φ and U, are random. Sometimes, a slightly different model is proposed where φ and U are still unobserved, but fixed, non-random quantities. For instance, under this setting (1) becomes

$$(X_{ij})_{i<j} \sim \bigotimes_{i<j} \mathrm{Be}(M_{\varphi(i)\varphi(j)}),$$

for a given, unknown, $\varphi : \{1, \dots, n\} \to \{1, \dots, k\}$, and the data distribution is denoted $P_{\varphi,M}$. Such models will be referred to as *fixed design* SBM and *random design* SBM respectively. Some theoretical arguments simplify in the fixed design case, for which the data distribution is a product measure, rather than a mixture of products measures. Most results below are obtained in both fixed and random design case.

Estimation in SBMs. We briefly review some existing results on SBMs. For parameters that can be expressed as moments, such as the edge density central limit theorems can be derived for empirical moments, see [8], [2].

Consider the estimation of the parameters π and M of the SBM itself, for a fixed (and known) number of classes k. The labels, given by the mapping φ, are not observed. If they were observed, Lemma 2 in [7] shows that the model is locally asymptotically normal with rescaling rates respectively $1/\sqrt{n}$ and $1/n$ at any 'interior' point (π_0, M_0) (see below). For known labels, this suggests a difference in the speed of estimation for the proportions vector π and connectivity matrix M. Remarkably, Theorems 1-2 of [7] show that in the original model (1)—that is, for unobserved labels—the parameters π and M can be estimated at rate respectively $1/\sqrt{n}$ and $1/n$. Denote by $\hat{\pi}$, $\hat{M}$ the profile maximum likelihood estimators of π, M. Then, as $n \to \infty$,

$$\begin{aligned} \sqrt{n}(\hat{\pi}^* - \pi_0) &\to \mathcal{N}(0, T_1) \\ n(\hat{M}^* - M_0) &\to \mathcal{N}(0, T_2) \end{aligned} \tag{4}$$

where $(\hat{\pi}^*, \hat{M}^*)$ denote some label switched-version of $(\hat{\pi}, \hat{M})$, which may be random and depend on n, and the convergence is in distribution, under P_{π_0,M_0}, to normal limits with covariance matrices T_1, T_2. It can also be of interest to consider sparser versions of the model, and [7] more generally obtain results when the connectivity matrix is normalised by a factor ρ that may go to 0 with n. For simplicity of exposition in the present paper we focus on the dense case, but results in a similar spirit could be derived for $\rho \to 0$.

The asymptotic result (4) holds under the assumption that the number of classes k is known and fixed as $n \to \infty$, say $k = k_0$, and that the true parameter (π_0, M_0) is such that all coordinates of π_0 are nonzero, and no two lines of the matrix M_0 coincide. This can be thought of as saying that the true parameter should be in the 'interior' of the parameter set. In the sequel, as we shall work with the *quadratic* risk, the rates in (4) will be squared, and we refer to a rate as *fast* if it is of order $1/n^2$, that is scaling as the number of edges, and as *slow* if it is of order $1/n$, that is scaling as the number of nodes.

Our contribution. The convergence (4) is asymptotic and pointwise at the interior point π_0, M_0. A natural question is whether anything changes if one aims at a rate result that would be non-asymptotic and uniform over the parameter set. To fix ideas, consider the following simple example. Take $n = 30$, $k = 5$, let π be fixed for simplicity to equiproportions $\pi = [1/5, 1/5, 1/5, 1/5, 1/5]$, and let the true connectivity matrix $M = M_0$ be given below. One also defines a matrix $\tilde{M}$ as

follows

$$M_0 = \begin{bmatrix} .55 & .45 & .4 & .1 & .7 \\ .45 & .55 & .4 & .1 & .7 \\ .4 & .4 & .6 & .2 & .1 \\ .1 & .1 & .2 & .1 & .4 \\ .7 & .7 & .4 & .4 & .2 \end{bmatrix}, \qquad \tilde{M} = \begin{bmatrix} .5 & .5 & .4 & .1 & .7 \\ .5 & .5 & .4 & .1 & .7 \\ .4 & .4 & .6 & .2 & .1 \\ .1 & .1 & .2 & .1 & .4 \\ .7 & .7 & .4 & .4 & .2 \end{bmatrix}.$$

One notes that the upper-left corner of M_0 is close to $\begin{bmatrix} .5 & .5 \\ .5 & .5 \end{bmatrix}$. Replacing the upper-left corner of M_0 by this matrix leads to the matrix $\tilde{M}$. Note that the first two lines of $\tilde{M}$ are equal and that it in fact corresponds to a SBM with $k = 4$ classes instead of 5. Does the fact that M_0 is close to $\tilde{M}$ change something in terms of estimation rates? To address this question, we consider the simplest possible statistical setting of SBMs where everything is known except one parameter, and we derive estimation rates for it. In this paper, we will focus on the *connectivity* parameters, that is on the elements of M, so the proportions vector π will typically be kept fixed. Section 2 considers the toy example where $k = 2$. It will be seen that the estimation rate drops if one requires uniformity. In Section 3, we consider SBMs with k classes. We obtain a minimax lower bound for estimation of a parameter of the connectivity matrix that generalises the case $k = 2$. The number k of classes is allowed to depend on n, and the result is *local* around *any* possible SBM. For instance, it can be applied to the example with M and $\tilde{M}$ above. We also show that the rate is sharp for 'most' connectivity matrices M, at least if k does not grow too fast with n. In Section 4, we derive some results for certain functionals of the graphon model (2) and conclude with a short discussion. Appendices A-D collect proofs and a number of useful lemmas.

Notation. Denote by $[k]^n$ the set of all mappings $\{1, \dots, n\} \to \{1, \dots, k\}$. A SBM is entirely specified through three parameters: the number of classes k, a vector of proportions π that belongs to the k-dimensional simplex, and the connectivity matrix M of size $k \times k$. The representation of the resulting distribution $P_{\pi,M}$ as a mixture is useful to make the link between the standard random design case and the fixed design case. We have

$$P_{\pi,M} = \sum_{\varphi \in [k]^n} \mu_\pi[\varphi] \bigotimes_{i<j} \mathrm{Be}(M_{\varphi(i)\varphi(j)}), \tag{5}$$

where, for $\mathcal{M}(n; \pi_1, \dots, \pi_k)$ the multinomial distribution with parameters $n, \pi_1, \dots, \pi_k$ and $N_j(\varphi) = \sum_{i=1}^n \mathbb{1}_{\varphi(i)=j}$ the number of times the label j is present, the multinomial probabilities are given by

$$\mu_\pi[\varphi] = Pr\left[\mathcal{M}(n; \pi_1, \dots, \pi_k) = (N_1(\varphi), \dots, N_k(\varphi))\right].$$

In the fixed design model, the labels given through φ are fixed and unobserved, so that the distribution is $P_{\varphi,M}$ given by

$$P_{\varphi,M} = \bigotimes_{i<j} \mathrm{Be}(M_{\varphi(i)\varphi(j)}).$$

For A a subset of the integers, let $|A|$ denote its cardinality. Let $ER(p)$ denote the Erdös-Renyi distribution with parameter p over n nodes, that is $ER(p) = ER(p, n) = \otimes_{i<j}\mathrm{Be}(p)$.

2. Toy example: the case $k = 2$

In this section, we consider the case $k = 2$. Let, for $\theta \in [-1/2, 1/2]$,

$$e = \left[\frac{1}{2}, \frac{1}{2}\right] \tag{6}$$

$$Q^\theta = \begin{bmatrix} \frac{1}{2} + \theta & \frac{1}{2} - \theta \\ \frac{1}{2} - \theta & \frac{1}{2} + \theta \end{bmatrix}. \tag{7}$$

Recalling that $P_{\pi,M}$ denotes the distribution of the data X in the SBM (1) with parameters π, M, consider the submodel

$$\mathcal{M} = \left\{P_\theta := P_{e,Q^\theta}, \quad \theta \in [-1/2, 1/2]\right\}. \tag{8}$$

The set $\mathcal{M}$ is a 1–dimensional submodel through the set of all SBMs with at most two classes. For $\theta = 0$ the matrix Q^0 is degenerate and the model is simply an Erdös-Reyni graph model with parameter $1/2$, that is all edges are independent and have a probability $1/2$ of being present. In this notation, the case $\theta = 0$ is somewhat overspecified and could also more simply be given by 1-dimensional vectors $\tilde{e} = [1]$, $\tilde{q} = [1/2]$ (note that the fact that there are various parameterisations for the case $\theta = 0$ causes no identifiability issue, as the probability measure P_0 is well identified and distinct from all P_θ, $\theta \neq 0$). Connectivity matrices with two parameters, one for intra-group and one for between-group connections (so-called affiliation models), are common in the literature on random graphs models, see e.g. [2, 4, 25] and references therein.

In the fixed design case, the model is $\mathcal{M}_F = \{P_{\theta,\varphi} := P_{\varphi,Q^\theta}, \quad \theta \in [-1/2, 1/2],\ \varphi \in [2]^n\}$. The expectations with respect to the measures P_θ and $P_{\theta,\varphi}$ are denoted respectively E_θ and $E_{\theta,\varphi}$.

Theorem 1. *Consider a stochastic blockmodel (1) with $k = 2$ specified by $\mathcal{M}$, that is $P_\theta = P_{e,Q^\theta}$ with e, Q^θ given by (6)-(7). There exists a constant $c_1 > 0$ such that for all $n \geq 2$,*

$$\inf_T \sup_{\theta \in [-1/2,1/2]} E_\theta \left[T(X) - \theta\right]^2 \geq \frac{c_1}{n},$$

where the infimum is taken over all estimators T of θ in the model $\mathcal{M}$.

Theorem 1 states that, even in a very simple SBM with $k = 2$ classes and only one unknown parameter in its connectivity matrix, the minimax estimation rate is no faster than $1/n$. This is no contradiction to the fast rate obtained in [7] (meaning a $1/n$ rate for the convergence in distribution but a $1/n^2$ rate for the quadratic risk): the latter is a pointwise asymptotic result, and assumes that no two lines of the connectivity matrix are the same, whereas Theorem 1 is nonasymptotic and uniform. It shows that the rate in a two-class model changes for distributions close to an Erdős-Renyi model ($k = 1$); informally, models close to the 'boundary' are harder to estimate. We note the result does not require the sub-model $\mathcal{M}$ to include the Erdős-Renyi model; see the remark below. The phenomenon is reminiscent of effects familiar from community detection, where matrices similar to (7) naturally arise as most difficult submodels. Community detection is a *testing* problem, though, as opposed to the *estimation* problem considered here.

Remarks. (a) In Theorem 1, one can take $c_1 = 1/107$; additionally, the supremum can be restricted to $(-\theta_n, \theta_n)$ for

$$\theta_n = \frac{c_0}{\sqrt{n}} \qquad \text{and} \quad c_0 = \frac{1}{3 \cdot 2^{3/4}} \approx 0.56.$$

Moreover, the proof implies that one can restrict the supremum to a set not actually containing $\theta = 0$, but rather two points close enough to $\theta = 0$, namely $\theta_1 = c_1/\sqrt{n}$, $\theta_2 = c_2/\sqrt{n}$ for suitably chosen, fixed constants $c_1, c_2 > 0$.

(b) A similar result as in Theorem 1 holds in the fixed design model, to wit

$$\inf_{T_f} \sup_{\theta \in [-1/2,1/2],\, \varphi \in [2]^n} E_{\theta,\varphi} \left[T_f(X) - \theta\right]^2 \geq \frac{c_1}{n},$$

where the infimum is taken over all estimators T_f of θ in the fixed design model.

Theorem 1 provides a lower bound. There is a corresponding, matching upper-bound, which we obtain next, by defining an estimator of θ whose maximum quadratic risk matches the lower bound

of Theorem 1. To do so, let, for any σ an element of $[2]^n$, that is a mapping $\{1,\dots,n\} \to \{1,2\}$,

$$2Z_n(\sigma, X) := -\sum_{i<j,\ \sigma(i)=\sigma(j)} (1-2X_{ij}) + \sum_{i<j,\ \sigma(i)\neq\sigma(j)} (1-2X_{ij}). \tag{9}$$

Maximising (9) in σ leads to set

$$\hat{\sigma} = \operatorname*{argmax}_{\sigma\in[2]^n} |Z_n(\sigma, X)|$$

which leads to the profile maximum likelihood estimate

$$\hat{\theta} = \frac{Z_n(\hat{\sigma}, X)}{b_n} \qquad \text{and} \qquad b_n = \binom{n}{2} = \frac{n(n-1)}{2}. \tag{10}$$

This estimator can be used both in the fixed and random design case and can be interpreted as a (pseudo)-maximum likelihood estimate, see Appendix B.

Theorem 2. *Consider a stochastic blockmodel (1) with $k = 2$ specified by $\mathcal{M}$, that is, $P_\theta = P_{e,Q^\theta}$ with e, Q^θ given by (6)-(7). Let $\hat{\theta}$ be the estimator defined by (10). There exists a constant $C_1 > 0$ such that for all $n \geq 2$,*

$$\sup_{\theta\in[-1/2,1/2]} E_\theta \left[\hat{\theta} - \theta\right]^2 \leq \frac{C_1}{n}.$$

The same risk bound holds for $\hat{\theta}$ in the fixed design model, uniformly over θ and $\varphi \in [2]^n$.

Remark (different parameter choices). It is not hard to check that the results of Theorems 1-2 are unchanged if instead of $1/2$ in the matrix Q^θ in (7), another number $a_0 \in (0,1)$ is used. If a_0 is bounded away from 0 and 1, assuming $\min(a_0, 1-a_0) \geq \rho > 0$, then the results are only modified by constants. Also, if the proportions vector π is of the form $[b\,,\,1-b]$ with $b > 0$, similar results continue to hold, provided the matrix Q^θ is replaced by

$$Q_b^\theta = \begin{bmatrix} \frac{1}{2} + c_b\theta & \frac{1}{2} - c_b\theta \\ \frac{1}{2} - c_b\theta & \frac{1}{2} + c_b\theta \end{bmatrix}$$

for a suitable c_b that depends on b.

Remark (related submodels). Another natural submodel to consider is

$$\check{Q} = \begin{bmatrix} \frac{1}{2} + \theta & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}. \tag{11}$$

Results as in Theorems 1-2, with the same rates, can be derived for this submodel as well. As its form suggests, proving the upper bound is easier (for given π the edge density is an affine function of θ, so a first moment estimator works), whereas the lower bound is somewhat more involved. We prefer to work with Q^θ in (7), as we are particularly interested in the behaviour of MLE-type estimators, see also Section 3.

Theorems 1 and 2 combined show the quadratic rate of estimation along the submodel $\mathcal{M}$ for estimating a connectivity parameter is exactly of order n^{-1}, up to constants. This 'slow' rate (as compared to n^{-2}) arises even if all other parameters—here, the vector of proportions π—are assumed known. The submodel built for $k = 2$ can be regarded as a local perturbation of an Erdös-Renyi graph model with connection probability $1/2$. The drop in the rate is already noteworthy, as the rate of estimation of p for a ER(p) model is of the order n^{-2}. More generally, however, one may consider perturbation around a SBM with $k-1$ classes instead of the ER(p) model, and ask whether similar results hold for other choices of k, and whether the rate changes with k. More precisely: Can one describe convergence rates along simple submodels, in one unknown connectivity parameter, that still exhibit slower rates of convergence than the expected pointwise rate?

3. Main result and local minimax bounds for a SBM with k classes

To extend our results for the case $k = 2$, we note the connectivity matrix of a SBM with at most $k-1$ classes is of the form

$$M = \begin{bmatrix} a_0 & a_1 & \cdots & a_{k-2} \\ a_1 & b_{11} & \cdots & b_{1k-2} \\ \vdots & \vdots & & \vdots \\ a_{k-2} & b_{1k-2} & \cdots & b_{k-2k-2} \end{bmatrix} \qquad \text{where } a_0, a_i, b_{ij} \in [0,1] \text{ for } i,j \in [k-2]\,. \tag{12}$$

One can ensure the number of classes is exactly $k-1$ by requiring no two rows of M coincide, but this will not be needed in the sequel. For simplicity of notation, and comparability to Section 2, we assume $a_0 = 1/2$ throughout. Results are easily adapted to the case $a_0 \in (0,1)$, requiring only that a_0 be bounded away from 0 and 1, as in the remark above.

We consider 1-dimensional submodels in the parameter space of connectivity matrices: Set

$$e_k = \left[\frac{1}{k}, \cdots, \frac{1}{k}\right], \tag{13}$$

and, for coefficients $\{a_i\}, \{b_{ij}\}$ as above, define

$$M^\theta = \begin{bmatrix} \frac{1}{2}+\theta & \frac{1}{2}-\theta & a_1 & \cdots & a_{k-2} \\ \frac{1}{2}-\theta & \frac{1}{2}+\theta & a_1 & \cdots & a_{k-2} \\ a_1 & a_1 & b_{11} & \cdots & b_{1k-2} \\ \vdots & \vdots & \vdots & & \vdots \\ a_{k-2} & a_{k-2} & b_{1k-2} & \cdots & b_{k-2k-2} \end{bmatrix} = \begin{bmatrix} Q^\theta & A \\ A^T & B \end{bmatrix}, \tag{14}$$

where

$$Q^\theta = \begin{bmatrix} \frac{1}{2}+\theta & \frac{1}{2}-\theta \\ \frac{1}{2}-\theta & \frac{1}{2}+\theta \end{bmatrix}, \quad A = \begin{bmatrix} a_1 & a_2 & \ldots & a_{k-2} \\ a_1 & a_2 & \ldots & a_{k-2} \end{bmatrix}, \quad B = \begin{bmatrix} b_{11} & \cdots & b_{1k-2} \\ \vdots & & \vdots \\ b_{1k-2} & \cdots & b_{k-2k-2} \end{bmatrix}.$$

Thus, M^θ is a symmetric $k \times k$ matrix, obtained from M by replacing the scalar coefficient a_0 by the 2×2 matrix Q^θ, and repeating the vector $(a_i)_{1 \le i \le k-2}$. For simplicity, we choose π in (13) equiproportional, although (as in the case $k = 2$) analogous results can be obtained if the proportions are of similar sizes.

Now consider the model defined by, for e_k, M^θ as in (13)-(14),

$$\mathcal{M}_k = \left\{P_\theta := P_{e_k, M^\theta}, \quad \theta \in [-1/2, 1/2]\right\}. \tag{15}$$

The set $\mathcal{M}_k$ is a 1–dimensional submodel through the set of all SBMs with at most k classes. As before for $\theta = 0$ the matrix M^0 has two identical rows, and the model becomes a SBM with at most $k-1$ classes, with connectivity matrix given by M defined above. Note that the model is in a sense quite 'simple', as only θ varies, all other parameters being fixed and known. Theorem 3 below shows that estimation of θ in this model is in fact fairly slow.

3.1. Lower bound result

As before let us denote by E_θ the expectation under P_θ in the model $\mathcal{M}_k$ given by (15).

Theorem 3. *Consider a stochastic blockmodel (1) with $k \ge 2$ classes specified by $\mathcal{M}_k$ in (15), that is $P_\theta = P_{e_k, M^\theta}$ with e_k, M^θ given by (13)-(14), for fixed matrices A, B with arbitrary coefficients. There exists a constant $c_3 = c_3(\rho) > 0$, independent of A, B, such that, for all $n \ge 12k$,*

$$\inf_T \sup_{\theta \in [-1/2, 1/2]} E_\theta \left[T(X) - \theta\right]^2 \ge c_3 \frac{k}{n},$$

where the infimum is taken over all estimators T of θ in the model $\mathcal{M}_k$.

Theorem 3 establishes that the minimax estimation rate of θ in the model (15) is at best of the order k/n, uniformly over k and n. An intuitive explanation for this particularly slow rate is as follows: the phenomenon observed for $k = 2$ is still present but this time the part of the matrix Z^θ containing information about θ is smaller, as only of the order $2/k$ of the nodes will be assigned to classes 1 or 2, which are the elements of the connectivity matrix that depend on θ.

An important point is that this lower bound is minimax *local* (as opposed to more commonly proved minimax global results) that is, not only does this slow rate occur around one specific least-favorable point in the parameter space, it does occur around *any* point. More precisely, for any $k \geq 2$ and the proportions vector being fixed, whatever connectivity matrix M with $k-1$ classes as in (12) we start from, there exists at least one submodel around M, namely $\mathcal{M}_k$ in (15), such that estimation of a connectivity parameter in M cannot be faster than k/n. One could have wondered from Theorem 1 if the slow rate there was not due to the fact that the model for $\theta = 0$ was quite special, namely of Erdös-Reyni type. This is not the case. Proving such a local bound makes the proof of Theorem 3 particularly non-trivial, as one has to quantify the L^1-distance between two mixtures of probability measures, instead of between one fixed measure and a mixture.

It is also of interest to compare the rate in Theorem 3 to the one that would be obtained if the labels were observed. Taking for simplicity an asymptotic setting with $n \to \infty$ and $k/n \to 0$, an adaptation of Lemma 2 in [7] leads to a squared rate in $(k/n)^2$ (extending the *quadratic* rate n^{-2} for the connectivity parameters in (4) corresponding to the fixed k case considered in [7]). The rate in Theorem 3 is the square-root of this rate and thus much slower.

3.2. Upper bound result

To obtain a corresponding upper bound, we consider a specific estimator in the model $\mathcal{M}_k$, which is similar to the profile likelihood approach in Section 2, with some adaptations. Let Σ_e denote the set of all maps $\sigma \in [k]^n$ such that, for some constants c_1, c_2, for any $1 \leq j \leq k$,

$$c_1 \frac{n}{k} \leq |\sigma^{-1}(j)| \leq c_2 \frac{n}{k}.$$

The set Σ_e thus consists of those maps σ that produce k classes of roughly similar size n/k. For a given subset $S \subset \{1, 2, \ldots, n\}$, define

$$2Z_n(\sigma, S, X) = - \sum_{i<j,\ i,j\in S,\, \sigma(i)=\sigma(j)} (1 - 2X_{ij}) + \sum_{i<j,\ i,j\in S,\, \sigma(i)\neq\sigma(j)} (1 - 2X_{ij}). \tag{16}$$

This criterion function is similar to the one in (9), but restricts to a subset S of indices. The set $S_I := \varphi^{-1}(\{1, 2\})$ of vertices with labels 1 or 2 is not observed. We note that, if S_I was known, one could follow an argument similar to the one leading to (10): Obtain an estimate $\hat{\sigma}_I$ of $\sigma_{|\{1,2\}}$ by maximising $Z_n(\sigma, S_I, X)$, and then set $\bar{\theta} := Z_n(\hat{\sigma}_I, S_I, X)/n_k$ for $n_k = \binom{\lfloor n/k \rfloor}{2}$. It can then be checked that $\bar{\theta}$ has a quadratic risk uniformly at most of the order k/n, which follows relatively easily as for Theorem 2. To obtain the result below, where the labels are not observed, we estimate the unknown set $\varphi^{-1}(\{1, 2\})$ from data.

We can simplify the proof of Theorem 4 below, without loss of generality, by maximizing over a grid. Let $\Theta_n = \{i/(2n^2),\ i = -n^2, \ldots, n^2\}$ be a regular grid within $\Theta = [-1/2, 1/2]$, and define

$$(\tilde{\sigma}, \tilde{\theta}) := \operatorname*{argmin}_{\sigma\in\Sigma_e,\, \theta\in\Theta_n} \sum_{i<j} (X_{ij} - Z^\theta_{\sigma(i)\sigma(j)})^2 \tag{17}$$

$$\tilde{S}_I := \tilde{\sigma}^{-1}(\{1, 2\}). \tag{18}$$

Equation (17) defines a global maximum-likelihood type estimator, which is then used to obtain an estimate $\tilde{S}_I$ of the set of nodes labelled 1 or 2. Given this estimate, one can apply the profile-type method already used in the case $k = 2$: For $n_k = \binom{\lfloor n/k \rfloor}{2}$, $\tilde{S}_I$ as in (18), and Z_n as in (16), set

$$\hat{\sigma}_I = \operatorname*{argmax}_{\sigma \in \Sigma_e} |Z_n(\sigma, \tilde{S}_I, X)| \tag{19}$$

$$\hat{\theta} = \frac{Z_n(\hat{\sigma}_I, \tilde{S}_I, X)}{n_k}. \tag{20}$$

We require the coefficient a_0 of the matrix M in (12) to be sufficiently distinct from the remaining entries: Let $\mathcal{C} = \{a_i, b_{ij}, 1 \leq i, j \leq k-2\}$ be the set of coefficients of the matrices A and B in (14), with $a_0 = 1/2$,

$$\min_{c \in \mathcal{C}} \{|c - a_0|\} \geq 2\kappa > 0. \tag{21}$$

Theorem 4. *Consider a stochastic blockmodel (1) with $k \geq 2$ classes specified by $P_\theta = P_{e_k, Z^\theta}$ with e_k, Z^θ given by (13)-(14), for fixed matrices A, B. Define $\hat{\theta} = \hat{\theta}(X)$ as in (20). Suppose (21) holds and that, for some small enough d and κ as in (21),*

$$k^3 \log k \leq d\kappa^4 n. \tag{22}$$

Then there exists a universal constant $C_1 > 0$ such that for $n \geq 5$,

$$\sup_{|\theta| \leq \kappa} E_\theta \left[\hat{\theta} - \theta\right]^2 \leq C_1 \frac{k}{n}.$$

Also, the same risk bound holds for $\hat{\theta}$ in the fixed design model, uniformly over $|\theta| \leq \kappa$ and $\varphi \in \Sigma_e$.

Note κ in (22) may depend on k and n, and may go to zero in a framework where k, n go to infinity. Two examples for the behaviour of κ are given below. A condition such as (21) for the procedure (20) is not unexpected. Indeed, it is based on a preliminary estimate of the set of nodes with labels 1 or 2, and may not be effective if this estimation step is not sufficiently accurate, or not possible. We discuss the necessity of a condition on M that ensures the rate k/n is attainable below. On the other hand, when k is not too large with respect to n, 'most' matrices will satisfy the condition (22) as we show in the second example below. The rate k/n is then typically achievable, at least locally in a neighborhood of size κ around 0.

Example 1 (well-separated block). If κ is a fixed positive constant e.g. $1/4$, then the submatrix Q^θ is well separated from the other coefficients of the matrix M. The procedure above then correctly picks up a sensible approximation of the true set $\sigma^{-1}(\{1, 2\})$ via $\tilde{S}_I$ and the rate k/n is achieved, as long as k does not grow faster than $n^{1/3}/\log n$, an already fairly important number of classes.

Example 2 (randomly sampled matrix M). Suppose that the symmetric matrix $M =: (c_{ij})$ in (12) is a random matrix whose upper triangular entries are drawn i.i.d. with uniform distribution $\mathcal{U}[0, 1]$, except $c_{11} = 1/2$. The distribution of $|c_{ij} - 1/2|$ except for $i = j = 1$ is then $\mathcal{U}[0, 1/2]$, and it is a standard fact that the first order statistic of an uniformly distributed sample of size N is $\text{Beta}(1, N)$ distributed. That implies the random variable $2 \min_{c_{ij} \in \mathcal{C}} |c_{ij} - 1/2|$ has law $\text{Beta}(1, k(k-1)/2 - 1)$. Therefore, κ in (21) is of order no less than $1/k^2$ with high probability. From (22) one deduces that for k of the form n^δ with $\delta < 1/11$ and n large enough, the rate k/n is achieved uniformly and locally, for typical matrices M. Inspection of the proof of Theorem 4 reveals that $k = o(n^\delta)$ with $\delta < 1/7$ in fact suffices for the rate k/n to be attained with high probability when M is random: this is achieved by distinguishing c_{ij} of the types a_i or b_{ij} in the proof and noting that the minimum of $|a_i - 1/2|$ over i will be of larger order k^{-1}, instead of k^{-2} for the minimum over i, j of $|b_{ij} - 1/2|$.

3.3. Necessity of conditions on M

The previous subsection shows that the rate k/n is achieved under a condition on M and/or k. In general, we expect the rate to depend on the matrices M. Although we do not investigate this point in full here, we discuss it briefly.

The estimation method above is applicable if connectivity parameters of the upper-left 2×2 block of M are separated sufficiently from the remaining entries. Among those matrices M^θ whose upper-left corner equals Q^θ, a worst case for this method should hence be a matrix whose coefficients in A and B all equal $1/2$. If specifically $Q^\theta = \check{Q}^\theta$ as in (11), this leads to a matrix

$$\check{M}^\theta = \begin{bmatrix} \frac{1}{2}+\theta & \frac{1}{2} & \frac{1}{2} & \cdots & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \cdots & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \cdots & \frac{1}{2} \\ \vdots & \vdots & \vdots & & \vdots \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \cdots & \frac{1}{2} \end{bmatrix}, \tag{23}$$

which is of course heavily over-specified from the SBM perspective. Consider the SBM in a fixed design case, where $\varphi : \{1, \dots, n\} \to \{1, \dots, k\}$ is unobserved. Suppose $\varphi^{-1}(1)$ is of size k/n, and the connectivity matrix is given by (23). This specific model can be regarded as a special case of the setting considered from a testing perspective in [10, 4]. From Theorem 2.2 in [10], one can deduce that the minimax rate for the quadratic risk when estimating θ is no better than $\rho_n = \max\left(\frac{k \log k}{n}, \frac{k^4}{n^2}\right)$, for $k, n \to \infty$ and $\rho_n = o(1)$. The rate is therefore no better than $(k \log k)/n$ even if k is small; it becomes much slower than k/n for large k. In the present paper, the upper-left 2×2 matrix in (23) is M^θ instead of $\check{M}^\theta$. For this case, we conjecture the obtained rate to be even slower. This will be investigated elsewhere.

4. Results for a class of functionals of smooth graphons

Stochastic block models can be identified with piece-wise constant graphons; we now consider the case where the graphon is a smooth function instead. Let $w : [0,1]^2 \to [0,1]$ a measurable function, let $\langle w \rangle$ be its graphon equivalence class, and denote by $P_{\langle w \rangle} = P_w$ the distribution of data X generated by the graphon model (2). Consider the problem of estimating the functional

$$\tau(\langle w \rangle) = \left[\int_{[0,1]^2} \left(w(x,y) - \int_{[0,1]^2} w \right)^2 dxdy \right]^{1/2}, \tag{24}$$

for any representer w of $\langle w \rangle$. This is well defined in terms of the graphon, as the integral is invariant under any simultaneous (Lebesgue-)measure-preserving transformation of x and y.

The statistic (24) can be interpreted as a 'graphon-standard deviation'. Its estimation under a smooth graphon model is, in a sense, analogous to the problem of estimating the functional θ in the simple SBM with two classes discussed in Section 2: Let w_θ be the piece-wise constant graphon characterizing the SBM defined by (6)-(7). Since $\tau(\langle w_\theta \rangle) = |\theta|$, estimating θ is then indeed equivalent to estimation of $\tau(\langle w \rangle)$ (if one considers only positive values of θ).

Under a 2-class SBM, the results of Section 2 show θ in (7) cannot be estimated faster than c/n. It is natural to ask whether the same still holds if one works with 'smoother' graphons instead of histograms (where we refer to $\langle w \rangle$ as smooth if at least one of its representers is a smooth function). The following result answers this question, both for τ and for a larger class of functionals containing τ.

Let $\mathcal{P}_B$ be the collection of all graphons that admit a representer which is a polynomial in x, y, with degree bounded by some integer $D \geq 2$ and coefficients bounded by an arbitrary constant $M > 0$ (this boundedness restriction is only to ensure a matching upper-bound in the next result). For any $0 \leq \theta \leq 1$, let us denote by w_θ the function from $[0,1]^2$ to $[0,1]$ given by

$$w_\theta(x,y) = \frac{1}{2} - \theta(x - \frac{1}{2})(y - \frac{1}{2}) \tag{25}$$

and let w_0 denote the constant function equal to $1/2$.

Theorem 5. *Let X be data from the graphon model (2). There exist constants $c_1, c_2 > 0$ such that, for $\tau(\cdot)$ defined by (24),*

$$\frac{c_1}{n} \leq \inf_{\hat{\tau}} \sup_{w \in \mathcal{P}_B} E_{P_w} \left[\hat{\tau}(X) - \tau(\langle w \rangle)\right]^2 \leq \frac{c_2}{n},$$

where the infimum is taken over all possible estimators of $\tau(\langle w \rangle)$ in model (2).

Let ψ be an arbitrary functional defined on graphon equivalence classes such that for some $c > 0$ and w_θ the function in (25),

$$|\psi(\langle w_\theta \rangle) - \psi(\langle w_0 \rangle)| \geq c|\theta|.$$

Then for some $d > 0$,

$$\inf_{\hat{\psi}} \sup_{w \in \mathcal{P}_B} E_{P_w} \left[\hat{\psi}(X) - \psi(\langle w \rangle)\right]^2 \geq \frac{d}{n}.$$

Thus, the minimax rate cannot be faster than c/n, even if one estimates a functional as elementary as (24), and restricts the model to a simple, small class of smooth graphons w—here, graphons with a polynomial representer. The functional

$$\psi(P_w) = \int_{[0,1]^2} \left| f(x,y) - \int_{[0,1]^2} |f(x,y)| dx dy \right| dx dy$$

is another example satisfying the assumptions of Theorem 5.

5. Discussion

In Gao et al. [16], the authors show that if one estimates the parameter function w of a graphon model, not observing the vertex labels—in this case, the variables U_i in (2)—does impact on the optimal rate. In the present paper, we have considered uniform estimation of certain functionals of graphon models (in particular, the loss function is quite different from theirs). For estimation of certain graph functionals—including the connectivity parameters considered in [7]—we have shown that the uniform, minimax rate does depend on whether the labels are observed, i.e. the phenomenon described in [16] persists even if one does not try to recover the entire function w. The fast quadratic rate $1/n^2$ is not achievable uniformly. If the number k of classes is known and fixed, the quadratic rate becomes $1/n$. If the number of classes k grows with n, the rate is much slower.

As already mentioned in the introduction, results in the same vein can be obtained in the "sparsified" case, where the function w is weighted by a rate function $\rho(n) \leq 1$ that decreases monotonically with n, as considered for example in [8, 30]; for the sake of conciseness, we have only considered the case $\rho = 1$.

Aspects of our proofs reflect the fact that graphon models constitute a specific type of mixture model, and estimation in mixtures can be difficult if mixture components are hard to distinguish; although no general theory of these phenomena seems to exist, we refer to the early work on estimation in finite mixture models by [19] and [6], and e.g. to [22] and [15] for more recent results.

Appendix A: Proofs of the lower bounds in SBMs

Notation. Recall that a SBM with k classes, proportions vector π and connectivity matrix M has distribution $P_{\pi,M}$ as given in (5). For A a $n \times n$ symmetric matrix with zero diagonal, we write

$$P_A = \bigotimes_{i<j} \mathrm{Be}(A_{i,j}).$$

If A is only given by $A_{i,j}$ for $i < j$, one extends it by symmetry and sets $A_{i,i} = 0$. This way, the distribution of a SBM in the fixed design case with given k, M and labelling function φ is P_{A^φ}, where $A^\varphi_{i,j} = A^\varphi(M)_{i,j} = M_{\varphi(i)\varphi(j)}$. In the random design case, if π is the vector with equal proportions $e_k = [k^{-1}, \dots, k^{-1}]$, then from (5),

$$P_{\pi,M} = P_{e_k,M} = \frac{1}{k^n} \sum_{\varphi \in [k]^n} P_{A^\varphi} \qquad \text{where} \qquad A^\varphi_{i,j} = M_{\varphi(i)\varphi(j)}.$$

In the sequel, C is a generic notation for a universal constant, whose value may change from line to line.

A.1. Two classes

The proof of Theorem 1 relies on Le Cam's method to bound the L_1 distance between a distribution and a finite mixture. This and other relevant results are summarized in Appendix D.

Proof of Theorem 1. Let $N = 2^n$ and let $A_1, \dots, A_N$ be the collection of symmetric matrices with general term $a_{ij} = a_{ij}(\theta, \varphi) = Q^\theta_{\varphi(i)\varphi(j)}$, $i < j$ and zero diagonal, for all possible $\varphi \in [2]^n$ and some $\theta \in \Theta$. Let A_0 be the matrix with all elements equal to $1/2$ on the off-diagonal, that is the matrix with $\theta = 0$. By Lemma 6, applied with $\tau = 0$ and $\theta = \theta_n$ small to be chosen below, in order to get a lower bound for the minimax risk, it is enough to bound the L^1-distance $\|\mathbb{P} - \mathbb{Q}\|_1$ between

$$\mathbb{P} = P_{A_0}, \quad \mathbb{Q} = \frac{1}{N} \sum_{k=1}^{N} P_{A_k}.$$

If $\lambda_1, \lambda_2, \mu$ are the probability distributions of the laws $\mathrm{Be}(q_1), \mathrm{Be}(q_2), \mathrm{Be}(1/2)$, a simple computation leads to

$$\int (\frac{d\lambda_1}{d\mu} - 1)(\frac{d\lambda_2}{d\mu} - 1) d\mu = (2q_1 - 1)(2q_2 - 1).$$

By Lemma 7 applied to $\mathbb{P}$ and $\mathbb{Q}$,

$$\begin{aligned} \|\mathbb{P} - \mathbb{Q}\|_1^2 &\le \frac{1}{4^n} \sum_{\varphi, \psi \in [2]^n} \prod_{i<j} (1 + \tau_{i,j}(\varphi, \psi)) - 1 \\ &\le \left[\frac{1}{4^n} \sum_{\varphi, \psi \in [2]^n} e^{\sum_{i<j} \tau_{i,j}(\varphi,\psi)} \right] - 1, \end{aligned}$$

where $\tau_{i,j}(\varphi, \psi) = (2m_{ij}(\varphi) - 1)(2m_{ij}(\psi) - 1)$. Note that $2m_{ij}(\varphi) - 1 = 2\theta(-1)^{\mathbb{1}_{\varphi(i) \neq \varphi(j)}}$ for any $\varphi \in [2]^n$. Denote $\eta_i = \mathbb{1}_{\varphi(i)=1} - \mathbb{1}_{\varphi(i)=2}$ and $\eta'_i = \mathbb{1}_{\psi(i)=1} - \mathbb{1}_{\psi(i)=2}$, for any index i. We have $2m_{ij}(\varphi) - 1 = 2\theta\eta_i\eta_j$, so that $\tau_{i,j}(\varphi, \psi) = 4\theta^2\eta_i\eta_j\eta'_i\eta'_j$. The term under brackets in the last display can be interpreted as an expectation over φ, ψ, where both variables are sampled uniformly from the set of all mappings from $\{1, \dots, n\}$ to $\{1, 2\}$. Under this distribution, the variables η_i for $i = 1, \dots, n$ are independent Rademacher, as well as the variables η'_i, and both samples are independent. Further note that the variables $R_i := \eta_i\eta'_i$ for $i = 1, \dots n$ form again a sample of independent Rademacher variables. It is thus enough to bound,

$$E\left[e^{4\theta_n^2 \sum_{i<j} R_i R_j} \right],$$

where E denotes expectation under the law of the R_i. The previous exponent is an instance of Rademacher chaos; its Laplace transform can be bounded using Lemma 4. If $Z_n := \sum_{i<j} R_i R_j$, we have that for any ε (say $\varepsilon = 1/2$), there exists $\lambda > 0$ such that for all $n \ge 2$,

$$E\left[e^{|Z_n|/(\lambda n)} \right] \le 1 + \varepsilon.$$

Choosing $n\theta_n^2 := 1/(4\lambda)$ leads to $\|\mathbb{P} - \mathbb{Q}\|_1 \leq \varepsilon = 1/2$, so that the minimax risk is bounded below by $(3/4)(16n\lambda)^{-1}$.

To obtain the constants as in the remark below the Theorem, using Lemma 5 in the final step of the proof with $\theta_n^2 = 1/(12s_n)$, $s_n^2 = n(n-1)/2$, $r(\cdot)$ as in Lemma 5, gives

$$R_M \geq \frac{\theta_n^2}{4}\left\{1 - \frac{1}{2}\sqrt{r(4\theta_n^2 s_n)}\right\},$$
$$\geq \frac{\sqrt{2}}{48n}\left\{1 - \frac{1}{2}\sqrt{r\left(\frac{1}{3}\right)}\right\} \geq \frac{0.45}{48n} \geq \frac{1}{107n}. \qquad \square$$

A.2. Lower bound for k classes

Here the problem is more delicate compared to $k = 2$ as the data distribution for $\theta = 0$, around which we build the lower bound, is itself a mixture (particularly in the random design case). As a first step, we establish a result in a fixed design setting, that is

$$\inf_T \sup_{\theta \in \Theta,\, \varphi \in [k]^n} E_{\theta,\varphi}\left[T - \theta\right]^2 \geq c\frac{k}{n} \qquad \text{for some } c > 0 \text{ and } k \geq 3\,. \tag{26}$$

Proof of (26). Define $m = m_k = 2\lfloor \frac{n}{k} \rfloor$. Set $S_1 = \{1, \ldots, m\}$ and $S_2 = \{m+1, \ldots, n\}$. Let $\varphi_0 \in [k]^n$ be a mapping such that

$$\varphi_0(S_1) \subset \{1, 2\} \qquad \text{and} \qquad \varphi_0(S_2) \subset \{3, \ldots, k\}. \tag{27}$$

Let $\varphi \in [k]^n$ be such that

$$\varphi(i) = \varphi_0(i) \quad \text{whenever } i \in S_2 \qquad \text{and} \qquad \varphi(S_1) \subset \{1, 2\}, \tag{28}$$

and denote by $\mathcal{F} = \mathcal{F}(\varphi_0, S_1)$ the set of all such φ's. Then the restriction $\varphi_{|S_1} =: \varphi_1$ of $\varphi \in \mathcal{F}$ to S_1 can be identified to an element of $[2]^m$.

Let M^θ be the $k \times k$ matrix defined in (14). For $\varphi \in \mathcal{F}$, let R_φ denote the matrix with general term $r_{ij} = r_{ij}(\varphi)$ equal to $M^\theta_{\varphi(i)\varphi(j)}$. There are as many such matrices as possible φ_1s, that is $|[2]^m| = 2^m$. As φ and φ_0 are identical by construction on S_2,

$$r_{ij} = \begin{cases} M^\theta_{\varphi_1(i)\varphi_1(j)} & \text{if } (i,j) \in S_1 \times S_1 \\ M^\theta_{\varphi_1(i)\varphi_0(j)} = a_{\varphi_0(j)-2} = M^0_{\varphi_0(i)\varphi_0(j)} & \text{if } (i,j) \in S_1 \times S_2 \\ M^\theta_{\varphi_0(i)\varphi_1(j)} = a_{\varphi_0(i)-2} = M^0_{\varphi_0(i)\varphi_0(j)} & \text{if } (i,j) \in S_2 \times S_1 \\ M^0_{\varphi_0(i)\varphi_0(j)} & \text{if } (i,j) \in S_2 \times S_2 \end{cases}$$

where φ_1 belongs to $[2]^m$. Next set, with A_0 the matrix with general term $m_{ij}(\varphi_0) = M^0_{\varphi_0(i)\varphi_0(j)}$,

$$\mathbb{P}' = P_{A_0} \qquad \text{and} \qquad \mathbb{Q}' = \frac{1}{2^m}\sum_{\varphi \in \mathcal{F}} P_{R_\varphi}.$$

Now we apply Lemma 7 to $\mathbb{P}', \mathbb{Q}'$. Both P_{A_0} and P_{R_φ} are product measures over all pairs of indices (i,j) with $1 \leq i < j \leq n$. By construction, the individual components of these products coincide as soon as either i or j does not belong to S_1. We write

$$P_{A_0} = \bigotimes_{i<j} P_{A_0}(i,j) \qquad \text{and} \qquad P_{R_\varphi} = \bigotimes_{i<j} P_{R_\varphi}(i,j).$$

where, for any indices i, j with $i < j$,

$$P_{A_0}(i,j) = \mathrm{Be}(M^0_{\varphi_0(i)\varphi_0(j)}) \qquad \text{and} \qquad P_{R_\varphi}(i,j) = \mathrm{Be}(M^\theta_{\varphi(i)\varphi(j)}).$$

For $\varphi, \psi \in [4]^n$, we set

$$\tau_{i,j}(\varphi,\psi) = P_{A_0}(i,j)\left[\left(\frac{dP_{R_\varphi}(i,j)}{dP_{A_0}(i,j)} - 1\right)\left(\frac{dP_{R_\psi}(i,j)}{dP_{A_0}(i,j)} - 1\right)\right].$$

If i or j belongs to S_2, then $r_{ij}(\varphi) = M^0_{\varphi_0(i)\varphi_0(j)} = A_0(i,j) = r_{ij}(\psi)$ by definition, in which case the last display equals 0. In Lemma 7, where φ, ψ play the role of the indices k, l. Identifying $\psi_{|S_1}$ with the corresponding mapping $\psi_1 \in [2]^m$, we have $\|\mathbb{P}' - \mathbb{Q}'\|_1^2 \leq \chi^2(\mathbb{Q}', \mathbb{P}')$ and

$$\chi^2(\mathbb{Q}', \mathbb{P}') \leq \frac{1}{2^{2m}} \sum_{\varphi_1,\psi_1 \in [2]^m} \prod_{1 \leq i < j \leq m} (1 + \tau_{i,j}(\varphi_1, \psi_1)) - 1. \tag{29}$$

The last expression coincides with the bound obtained in the proof of Theorem 1, with n replaced by $m = m_k$. As in that proof, there hence exist independent Rademacher variables $R_1, \ldots, R_m$ such that $Z_m = m^{-1}\sum_{1 \leq i < j \leq m} R_i R_j$ satisfies

$$\chi^2(\mathbb{Q}', \mathbb{P}') \leq \mathbb{E}\exp\left[4m\theta^2 |Z_m|\right] - 1.$$

Provided θ is defined as, for a a small enough constant,

$$\theta^2 = a(4m)^{-1} \sim \left(\frac{n}{k}\right)^{-1},$$

using Lemma 4 as in the proof of Theorem 1 leads to the bound $\|\mathbb{P}' - \mathbb{Q}'\|_1 \leq 1/2$ if θ^2 is a small enough multiple of k/n, which again leads to a lower bound for the minimax risk of a positive constant times k/n, which proves (26). □

Proof of Theorem 3. For $e = e_k$ and M^θ as in (13)-(14), let

$$\mathbb{Q}^\theta = \frac{1}{k^n} \sum_{\varphi \in [k]^n} \bigotimes_{i<j} \mathrm{Be}(M^\theta_{\varphi(i)\varphi(j)}),$$

and set $\mathbb{P} = \mathbb{Q}^0$ corresponding to $\theta = 0$. Our aim is to show that $\mathbb{Q}^\theta$ and $\mathbb{P}$ are close in the sense $\|\mathbb{Q}^\theta - \mathbb{P}\|_1 \leq 1/2$ say, while θ is a fixed positive multiple of $\sqrt{k/n}$. For a given $\varphi \in [k]^n$, set

$$S_1 := \varphi^{-1}(\{1,2\}) \qquad \text{and} \qquad S_2 := \varphi^{-1}(\{3,4,\ldots,k\}).$$

By definition we have $S_1 = S_2^c := \{1,\ldots,n\} \setminus S_2$ and $|S_1| + |S_2| = n$.

Step 1. One first shows that it is possible to restrict the sum in the definition of $\mathbb{Q}^\theta$ and $\mathbb{P}$ to φ's in the set

$$\begin{aligned}\mathcal{A}_n &= \left\{\varphi \in [k]^n, \quad \left||\varphi^{-1}(\{1,2\})| - \frac{2n}{k}\right| \leq \frac{n}{k}\right\} \\ &= \left\{\varphi \in [k]^n, \quad \left||\varphi^{-1}(\{3,\ldots,k\})| - \frac{(k-2)n}{k}\right| \leq \frac{n}{k}\right\}.\end{aligned}$$

The reason is that the large majority of sets S_1 have a cardinality of the order close to n/k. The proportion of φ's not in $\mathcal{A}_n$ among all possible φ's is given by the probability of a binomial $Y \sim \mathrm{Bin}(n, 2/k)$ variable being farther than n/k from its mean. By Bernstein's inequality, as $v := \mathrm{Var}[Y] = n(2/k)(1-2/k)$, for any $t > 0$,

$$\mathbb{P}\left[\left|Y - \frac{2n}{k}\right| > t\right] \leq 2\exp\left\{-\frac{t^2}{2v + t}\right\}.$$

Taking $t = n/k$ and setting $R_n := |\mathcal{A}_n|$, we have just shown that

$$0 \leq 1 - \frac{R_n}{k^n} \leq 2e^{-\frac{n}{k}(5 - \frac{8}{k})^{-1}}.$$

Now set

$$\widetilde{\mathbb{Q}}^\theta = \frac{1}{R_n} \sum_{\varphi \in \mathcal{A}_n} \bigotimes_{i<j} \mathrm{Be}(M^\theta_{\varphi(i)\varphi(j)}) \quad \text{and} \quad \widetilde{\mathbb{P}} = \widetilde{\mathbb{Q}}^0.$$

By the triangle inequality,

$$\|\mathbb{Q}^\theta - \mathbb{P}\|_1 \leq \|\mathbb{Q}^\theta - \widetilde{\mathbb{Q}}^\theta\|_1 + \|\widetilde{\mathbb{Q}}^\theta - \widetilde{\mathbb{P}}\|_1 + \|\widetilde{\mathbb{P}} - \mathbb{P}\|_1.$$

By Lemma 8, $\|\mathbb{Q}^\theta - \widetilde{\mathbb{Q}}^\theta\|_1 + \|\widetilde{\mathbb{P}} - \mathbb{P}\|_1$ is bounded above by $4(1 - R_n/k^n) \leq 8e^{-\frac{n}{k}(5-\frac{8}{k})^{-1}}$.

Step 2. We now focus on bounding the middle term $\|\widetilde{\mathbb{Q}}^\theta - \widetilde{\mathbb{P}}\|_1$. Let Σ_n denote the collection of subsets of $\{1, 2, \ldots, n\}$ with $|S_2 - (k-2)n/k| \leq n/k$. For a given $S \in \Sigma$, let $\varphi_S = \varphi_{|S}$ denote the restriction of φ to S. Below we use the notation $\sum_{\varphi_S}$ with the meaning that each term of the sum corresponds to a possible mapping φ_S, that is a given collection of values $(\varphi(i))_{i\in S} \in \{1, \ldots, k\}^S$.

To do so, we rewrite $\widetilde{\mathbb{Q}}^\theta$ and $\widetilde{\mathbb{P}}$ as 'mixtures of mixtures', by splitting the sum over φ into a sum over S_2, φ_{S_2} and φ_{S_1} given S_2. Specifying φ is equivalent to giving oneself S_2 (then $S_1 = S_2^c$), φ_{S_2} and $\varphi_{S_2^c} = \varphi_{S_1}$. Denote

$$P^\theta_\varphi(i,j) = \mathrm{Be}(M^\theta_{\varphi(i)\varphi(j)}) \qquad \text{and} \qquad P^\theta_\varphi = \bigotimes_{i<j} P^\theta_\varphi(i,j).$$

For given S_2 and φ_{S_2}, set

$$T^\theta_{\varphi,S_2} = \frac{1}{2^{n-|S_2|}} \sum_{\varphi_{S_1} \,|\, S_2, \varphi_{S_2}} P^\theta_\varphi,$$

where one sums over all possible mappings φ_{S_1}, while S_2 and φ_{S_2} are fixed. We have

$$\widetilde{\mathbb{Q}}^\theta = \sum_{S_2 \in \Sigma_n,\, \varphi_{S_2}} \frac{2^{n-|S_2|}}{R_n} T^\theta_{\varphi,S_2}.$$

Note that the above measures are normalised to be probability measures. Indeed, given $S_2 \in \Sigma_n$, there are $2^{|S_1|} = 2^{n-|S_2|}$ possible choices for φ_{S_1}. As $\widetilde{\mathbb{Q}}^\theta$ is of total mass one, we have

$$\sum_{S_2 \in \Sigma_n,\, \varphi_{S_2}} \lambda_{S_2} = 1 \qquad \text{for} \quad \lambda_{S_2} := \frac{2^{n-|S_2|}}{R_n}.$$

Using Pinsker's inequality $\|P - Q\|_1^2 \leq 2KL(P,Q)$, as well as the first part of Lemma 9 below,

$$\begin{aligned}
\|\widetilde{\mathbb{Q}}^\theta - \widetilde{\mathbb{P}}\|_1^2 &= \left\| \sum_{S_2 \in \Sigma_n,\, \varphi_{S_2}} \lambda_{S_2} T^\theta_{\varphi,S_2} - \sum_{S_2 \in \Sigma_n, \varphi_{S_2}} \lambda_{S_2} T^0_{\varphi,S_2} \right\|_1^2 \\
&\leq 2KL \left[\sum_{S_2 \in \Sigma_n,\, \varphi_{S_2}} \lambda_{S_2} T^\theta_{\varphi,S_2}, \sum_{S_2 \in \Sigma_n, \varphi_{S_2}} \lambda_{S_2} T^0_{\varphi,S_2} \right] \\
&\leq 2 \sum_{S_2 \in \Sigma_n,\, \varphi_{S_2}} \lambda_{S_2} KL\left[T^\theta_{\varphi,S_2}, T^0_{\varphi,S_2}\right] \\
&\leq 2 \max_{S_2 \in \Sigma_n,\, \varphi_{S_2}} KL\left[T^\theta_{\varphi,S_2}, T^0_{\varphi,S_2}\right].
\end{aligned}$$

It is now sufficient to uniformly bound the KL-divergence between. For simplicity, we denote

$$T^\theta_{\varphi,S_2} = \frac{1}{2^{n-|S_2|}} \sum_{\varphi_{S_1} \,|\, S_2} P^\theta_\varphi := \frac{1}{2^{n-|S_2|}} \sum_{\varphi_1} P^\theta_{\varphi_1,\varphi_2},$$

where $\varphi_2 = \varphi_{S_2}$ and $\varphi_1 = \varphi_{S_1}$, and φ is the pair (φ_1, φ_2). Set

$$\lambda_{S_1} = 2^{-(n-|S_2|)} = 2^{-|S_1|}.$$

Using the definition of T^{θ}_{φ,S_2} above as well as the second part of Lemma 9 below,

$$KL\left[T^{\theta}_{\varphi,S_2}, T^{0}_{\varphi,S_2}\right] = KL\left[\sum_{\varphi_1}\lambda_{S_1}P^{\theta}_{\varphi_1,\varphi_2}, \sum_{\varphi_1}\lambda_{S_1}P^{0}_{\varphi_1,\varphi_2}\right]$$
$$\leq \sum_{\varphi_1}\lambda_{S_1}KL\left[\sum_{\varphi_1'}\lambda_{S_1}P^{\theta}_{\varphi_1',\varphi_2}, P^{0}_{\varphi_1,\varphi_2}\right],$$

and further bounding the KL-divergence by the χ^2-divergence,

$$KL\left[T^{\theta}_{\varphi,S_2}, T^{0}_{\varphi,S_2}\right] \leq \sum_{\varphi_1'}\lambda_{S_1}\chi^2\left[\sum_{\varphi_1'}\lambda_{S_1}P^{\theta}_{\varphi_1',\varphi_2}, P^{0}_{\varphi_1,\varphi_2}\right]$$
$$\leq \max_{\varphi_1}\ \chi^2\left[\sum_{\varphi_1'}\lambda_{S_1}P^{\theta}_{\varphi_1',\varphi_2}, P^{0}_{\varphi_1,\varphi_2}\right].$$

Combining this with the previous bounds one deduces that

$$\|\mathbb{Q}^{\theta} - \mathbb{P}\|_1 \leq \|\widetilde{\mathbb{Q}}^{\theta} - \widetilde{\mathbb{P}}\|_1 + 8e^{-\frac{n}{k}(5-\frac{8}{k})^{-1}}$$
$$\leq \max_{S_2\in\Sigma_n,\varphi(S_2)}\ \max_{\varphi_1}\ \chi^2\left[\sum_{\varphi_1'}\lambda_{S_1}P^{\theta}_{\varphi_1',\varphi_2}, P^{0}_{\varphi_1,\varphi_2}\right] + 8e^{-\frac{n}{k}(5-\frac{8}{k})^{-1}}.$$

To conclude the proof, observe that the structure of the bound in the maximum in the last display is nearly identical to the quantities appearing in Equation (29) for the fixed-design case.

In the present case, we have a fixed mapping $\varphi : \{1,\dots,n\} \to \{1,\dots,k\}$, with $\varphi_1 = \varphi_{|S_1}$ and $\varphi_2 = \varphi_{|S_2}$, that plays the role of φ_0 in the fixed-design case. On the other hand, we have a collection of other mappings, say $\bar{\varphi}$, that coincide with φ on S_2, that is $\bar{\varphi}_2 = \bar{\varphi}_{|S_2} = \varphi_{|S_2} = \varphi_2$, and that cover all possible cases for the image of S_1, namely $\bar{\varphi}_1 = \bar{\varphi}_{|S_1} = \varphi_1'$. The only difference to the fixed-design case is that $|S_1|$ belongs to $[n/k, 3n/k]$, instead of being exactly $2\lfloor n/k\rfloor$, as specified in the definition of Σ_n above. That is, denoting as above $Z_m = m^{-1}\sum_{1\leq i<j\leq m} R_iR_j$, with $m_1 = |S_1|$,

$$\chi^2\left[\sum_{\varphi_1'}\lambda_{S_1}P^{\theta}_{\varphi_1',\varphi_2}, P^{0}_{\varphi_1,\varphi_2}\right] \leq \mathbb{E}\exp\left[4m_1\theta^2|Z_{m_1}|\right] - 1.$$

This bound is uniform over S_2, φ_2. As $m_1 \leq 3n/k$, if one chooses $\theta^2 \leq 1/(12\lambda n/k)$, with $\lambda = \lambda(1+\varepsilon)$ the constant in Lemma 5, then this Lemma implies that for any m_1 between n/k and $3n/k$, the chi-square distance in the last display is bounded by ε. Crucially, the constant λ in Lemma 4 is independent of the number of terms in the Rademacher chaos. Deduce

$$\|\mathbb{Q}^{\theta} - \mathbb{P}\|_1 \leq \varepsilon + 8e^{-\frac{n}{k}(5-\frac{8}{k})^{-1}}.$$

Choosing $n/k > 12$ makes this bound smaller than $\varepsilon + 4/5 < 1$ for $\varepsilon < 1/5$. □

Appendix B: Proofs of the upper bound results in SBMs

B.1. Two classes, pseudo-likelihood

We first justify the interpretation of the estimator $\hat{\theta}$ in (10) as a maximum (pseudo-)likelihood estimate. In the fixed design model, suppose the data is Gaussian $\mathcal{N}(\theta_{ij}, 1)$ instead of Bernoulli

Be(θ_{ij}). This suggests defining a (pseudo-)log-likelihood $\ell_n(\sigma,\theta)$ as follows, with $c_n = \binom{n}{2}\log(2\pi)$,

$$\begin{aligned}
-2\ell_n(\sigma,\theta) &= \sum_{i<j}(X_{ij} - Q^{\theta}_{\sigma(i)\sigma(j)})^2 + c_n \\
&= \sum_{i<j,\ \sigma(i)=\sigma(j)}(X_{ij} - (\frac{1}{2}+\theta))^2 + \sum_{i<j,\ \sigma(i)\neq\sigma(j)}(X_{ij} - (\frac{1}{2}-\theta))^2 + c_n \\
&= \binom{n}{2}\theta^2 + \theta\left(\sum_{i<j,\ \sigma(i)=\sigma(j)}(1-2X_{ij}) - \sum_{i<j,\ \sigma(i)\neq\sigma(j)}(1-2X_{ij})\right) + C_n(X),
\end{aligned}$$

for a constant $C_n(X)$ depending only on n and X. Setting $b_n = \binom{n}{2}$ and

$$2Z_n(\sigma,X) := -\sum_{i<j,\ \sigma(i)=\sigma(j)}(1-2X_{ij}) + \sum_{i<j,\ \sigma(i)\neq\sigma(j)}(1-2X_{ij}),$$

it is enough to study the function $g_n(\theta,\sigma) := b_n\theta^2 - 2Z_n(\sigma,X)\theta$, which satisfies

$$g_n(\theta,\sigma) = b_n\left(\theta - \frac{Z_n(\sigma,X)}{b_n}\right)^2 - \frac{Z_n(\sigma,X)^2}{b_n} \geq -\frac{Z_n(\sigma,X)^2}{b_n}.$$

Consequently, the pseudo maximum likelihood estimator $(\hat\theta,\hat\sigma)$ is given by (10) as claimed.

B.2. Upper bound result, two classes

Proof of Theorem 2. We first prove the result in the fixed design case. Let θ_0,σ_0 denote the true values of θ,φ. The aim is to show that $E_{\theta_0,\sigma_0}(\hat\theta-\theta_0)^2 \leq C/n$ holds uniformly in θ_0,σ_0. For a given $\sigma\in 2^{[n]}$,

$$r_{ij} := X_{ij} - 1/2 + (-1)^{\mathbb{1}_{\sigma_0(i)=\sigma_0(j)}}\theta_0 \qquad \text{and} \qquad R_n(\sigma) := \sum_{\sigma(i)=\sigma(j)} r_{ij} - \sum_{\sigma(i)\neq\sigma(j)} r_{ij}.$$

One can write, for any $\sigma\in 2^{[n]}$,

$$\begin{aligned}
Z_n(\sigma_0,X) &= b_n\theta_0 + \sum_{i<j}(-1)^{\mathbb{1}_{\sigma_0(i)\neq\sigma_0(j)}}r_{ij} = b_n\theta_0 + R_n(\sigma_0), \\
Z_n(\sigma,X) &= \theta_0\delta(\sigma,\sigma_0) + R_n(\sigma),
\end{aligned}$$

where we have set

$$\delta(\sigma,\sigma_0) = \sum_{i<j}(-1)^{\mathbb{1}_{\sigma_0(i)\neq\sigma_0(j)}}(-1)^{\mathbb{1}_{\sigma(i)\neq\sigma(j)}}. \tag{30}$$

For any $t>0$ and $t_n = M_2/\sqrt{n}$, and for a large enough M_2 to be chosen below,

$$\begin{aligned}
&P_{\theta_0}[\sqrt{n}|\hat\theta-\theta_0|\geq t] \\
&\quad= P_{\theta_0}[\sqrt{n}|\hat\theta-\theta_0|\geq t]\mathbb{1}_{|\theta_0|\leq t_n} + P_{\theta_0}[\sqrt{n}|\hat\theta-\theta_0|\geq t]\mathbb{1}_{|\theta_0|>t_n} \\
&\quad=: \qquad \mathcal{P}_1(t) \qquad + \qquad \mathcal{P}_2(t).
\end{aligned}$$

By definition of $\hat\theta$, with $\delta(\sigma,\sigma_0)$ defined in (30),

$$b_n\hat\theta = Z_n(\hat\sigma,X) = \theta_0\delta(\hat\sigma,\sigma_0) + R_n(\hat\sigma) \qquad \text{and} \qquad Z_n(\sigma_0,X) = \theta_0 b_n + R_n(\sigma_0).$$

For any $t \geq 4M_2$, using that $|\delta(\hat{\sigma}, \sigma_0)| \leq b_n$,

$$\begin{aligned}
\mathcal{P}_1(t) &\leq P_{\theta_0}\left[\sqrt{n}|\theta_0|\left|\frac{\delta(\hat{\sigma},\sigma_0)}{b_n} - 1\right| + \frac{\sqrt{n}}{b_n}|R_n(\hat{\sigma})| \geq t\right] \mathbb{1}_{|\theta_0|\leq t_n} \\
&\leq P_{\theta_0}\left[\frac{\sqrt{n}}{b_n}|R_n(\hat{\sigma})| \geq t - 2M_2\right] \\
&\leq P_{\theta_0}\left[\sup_{\sigma\in\Sigma}|R_n(\sigma)| \geq \frac{b_n}{\sqrt{n}}\frac{t}{2}\right].
\end{aligned}$$

For $\mathcal{P}_2(t)$, there are two cases, depending on the sign of θ_0,

$$\begin{aligned}
\mathcal{P}_2(t) &\leq P_{\theta_0}\left[\sqrt{n}|\hat{\theta}-\theta_0| \geq t\right]\mathbb{1}_{\theta_0>t_n} + P_{\theta_0}\left[\sqrt{n}|\hat{\theta}-\theta_0| \geq t\right]\mathbb{1}_{\theta_0<-t_n} \\
&\leq \left\{P_{\theta_0}\left[\sqrt{n}|\hat{\theta}-\theta_0| \geq t,\, Z_n(\hat{\sigma},X) \geq 0\right] + P_{\theta_0}\left[\sqrt{n}|\hat{\theta}-\theta_0| \geq t,\, Z_n(\hat{\sigma},X) < 0\right]\right\}\mathbb{1}_{\theta_0>t_n} \\
&\quad + \left\{P_{\theta_0}\left[\sqrt{n}|\hat{\theta}-\theta_0| \geq t,\, Z_n(\hat{\sigma},X) \geq 0\right] + P_{\theta_0}\left[\sqrt{n}|\hat{\theta}-\theta_0| \geq t,\, Z_n(\hat{\sigma},X) < 0\right]\right\}\mathbb{1}_{\theta_0<-t_n}
\end{aligned}$$

Let us discuss the term $\theta_0 > t_n$ first and note that if $Z_n(\hat{\sigma}, X) \geq 0$, then $Z_n(\hat{\sigma}, X) = |Z_n(\hat{\sigma}, X)| \geq |Z_n(\sigma_0, X)| \geq Z_n(\sigma_0, X)$ using the definition of $\hat{\sigma}$ as a maximum. First,

$$\begin{aligned}
&P_{\theta_0}\left[\sqrt{n}|\hat{\theta}-\theta_0| \geq t, Z_n(\hat{\sigma},X) \geq 0\right]\mathbb{1}_{\theta_0>t_n} \\
&\leq P_{\theta_0}\left[\frac{\sqrt{n}}{b_n}|Z_n(\hat{\sigma},X) - Z_n(\sigma_0,X) + R_n(\sigma_0)| \geq t,\, Z_n(\hat{\sigma},X) \geq 0\right]\mathbb{1}_{\theta_0>t_n} \\
&\leq P_{\theta_0}\left[\frac{\sqrt{n}}{b_n}(Z_n(\hat{\sigma},X) - Z_n(\sigma_0,X)) + \frac{\sqrt{n}}{b_n}|R_n(\sigma_0)| \geq t\right]\mathbb{1}_{\theta_0>t_n} \\
&\leq P_{\theta_0}\left[\frac{\sqrt{n}}{b_n}\theta_0(\delta(\hat{\sigma},\sigma_0) - b_n) + \frac{\sqrt{n}}{b_n}R_n(\hat{\sigma}) + \frac{\sqrt{n}}{b_n}|R_n(\sigma_0)| \geq t\right]\mathbb{1}_{\theta_0>t_n} \\
&\leq P_{\theta_0}\left[2\frac{\sqrt{n}}{b_n}\sup_{\sigma\in\Sigma}|R_n(\sigma)| \geq t\right],
\end{aligned}$$

where the last inequality uses $\delta(\hat{\sigma}, \sigma_0) - b_n \leq 0$ and $\theta_0 \geq 0$.

Second, as $Z_n(\hat{\sigma}, X) < 0$ implies $Z_n(\hat{\sigma}, X) < -|Z_n(\sigma_0, X)|$ by definition of the maximum,

$$\begin{aligned}
&P_{\theta_0}\left[\sqrt{n}|\hat{\theta}-\theta_0| \geq t,\, Z_n(\hat{\sigma},X) < 0\right]\mathbb{1}_{\theta_0>t_n} \\
&\leq P_{\theta_0}\left[Z_n(\hat{\sigma},X) < -|Z_n(\sigma_0,X)|\right]\mathbb{1}_{\theta_0>t_n} \\
&\leq P_{\theta_0}\left[Z_n(\hat{\sigma},X) < -|\theta_0|b_n + |R_n(\sigma_0)|\right]\mathbb{1}_{\theta_0>t_n} \\
&\leq P_{\theta_0}\left[\theta_0(\delta(\hat{\sigma},\sigma_0) + b_n) < 2\sup_{\sigma\in\Sigma}|R_n(\sigma)|\right]\mathbb{1}_{\theta_0>t_n} \\
&\leq P_{\theta_0}\left[2\sup_{\sigma\in\Sigma}|R_n(\sigma)| > \frac{t_n b_n}{8}\right],
\end{aligned}$$

where for the last inequality we have used the lower bound on δ obtained in Lemma 1.

The case $\theta_0 < -t_n$ is treated in a symmetric way, by distinguishing the two cases $Z_n(\hat{\sigma}, X) < 0$ and $Z_n(\hat{\sigma}, X) \geq 0$ respectively. To obtain a deviation bound for $\hat{\theta}$, it is enough to study the supremum of the process $|R_n(\sigma)|$. For any given σ and $y > 0$, by Hoeffding's inequality,

$$P_{\theta_0}[|R_n(\sigma)| > y] \leq 2\exp\{-2y^2/b_n\}.$$

A union bound now leads to

$$P\left[\sup_{\sigma\in\Sigma}|R_n(\sigma)| \geq y\right] \leq 2^n \exp\left\{-2y^2/b_n\right\} \qquad \text{for any } y > 0.$$

This bound is smaller than $\exp\{-2n\}$ if one chooses $y = n^{3/2}$. Combining the bounds obtained previously, and choosing M_2 above as $M_2 = 64$, one deduces

$$P[\sqrt{n}|\hat{\theta} - \theta_0| \geq t] \leq 6e^{-2n} \qquad \text{for any } t \geq 4.$$

The deviation bound in turn implies the bound in expectation

$$\begin{aligned} E[n(\hat{\theta} - \theta_0)^2] &= E[n(\hat{\theta} - \theta_0)^2 \mathbb{1}_{\sqrt{n}|\hat{\theta}-\theta_0|\leq 4}] + E[n(\hat{\theta} - \theta_0)^2 \mathbb{1}_{\sqrt{n}|\hat{\theta}-\theta_0|>4}], \\ &\leq 16 + 6ne^{-2n}, \end{aligned}$$

where for the second term we have used $|\hat{\theta} - \theta_0| \leq 1$, as Θ has diameter 1. This concludes the proof of Theorem 2 in the fixed design case.

In the random design case, one slightly updates the definition of r_{ij}. Here, the design is specified by φ, which is now random, but one can consider

$$\tilde{r}_{ij} = X_{ij} - E[X_{ij} \,|\, \varphi].$$

By definition, $E[\tilde{r}_{ij}] = E[E[\tilde{r}_{ij} \,|\, \varphi]] = 0$. Now one can follow the proof in the fixed design case by writing all statements conditionally on φ. As conditionally on φ the variables r_{ij} are independent and centered, the arguments leading to the various upper bounds remain unchanged. As the upper-bounds themselves do not depend on φ, the bounds also hold unconditionally. □

What remains to be shown is the bound on δ used above:

Lemma 1. *For any $\sigma_0, \sigma \in \Sigma$, with $\delta(\sigma, \sigma_0)$ defined in (30) and $b_n = \binom{n}{2}$,*

$$-\frac{7}{8} b_n \leq \delta(\sigma, \sigma_0) \leq b_n \qquad \text{for any } \; n \geq 5.$$

Proof. The upper bound corresponds to the number of terms in the sum. For the lower bound, denote $\mathcal{C}_1 = \sigma_0^{-1}(\{1\}) = \{i,\ \sigma_0(i) = 1\}$. By symmetry, one can always assume $|\mathcal{C}_1| \geq n/2$, otherwise one works with $\mathcal{C}_2 = \sigma_0^{-1}(\{2\})$. The number $T_\sigma(\mathcal{C}_1)$ of pairs $(i, j) \in \mathcal{C}_1 \times \mathcal{C}_1$ for which $\sigma(i) \neq \sigma(j)$ is at most $2N_1(\sigma)N_2(\sigma)$, if $N_i(\sigma) = |\sigma^{-1}(\{i\}) \cap \mathcal{C}_1|$, $i = 1, 2$. This implies $T_\sigma(\mathcal{C}_1) \leq |\mathcal{C}_1|^2/2$, using the inequality $p(q - p) \leq q^2/4$, for any $0 \leq p \leq q$. Thus the number of positive elements in the sum defining $\delta(\sigma, \sigma_0)$ is at least $(|\mathcal{C}_1|^2 - |\mathcal{C}_1|^2/2 - |\mathcal{C}_1|)/2$, where $|\mathcal{C}_1|$ corresponds to the diagonal terms and the division by 2 to the fact that the sum is restricted to $i < j$ only (note that the general term of the sum defining δ is symmetric in i, j). This is at least $|\mathcal{C}_1|^2/8$ if $|\mathcal{C}_1| \geq 3$. Hence,

$$\delta(\sigma, \sigma_0) \geq -b_n + 2\frac{|\mathcal{C}_1|^2}{8},$$

which is the desired bound in view of $|\mathcal{C}_1|^2 \geq b_n/2$. □

B.3. Upper bound result, k classes

Proof of Theorem 4. First consider the fixed design case: As in the proof of Theorem 2, let θ_0, σ_0 denote the true values of θ, φ. The aim is to show that

$$\sup_{|\theta_0| \leq \kappa,\ \sigma_0 \in \Sigma_e} E_{\theta_0, \sigma_0}(\hat{\theta} - \theta)^2 \leq C\frac{k}{n}.$$

Let us denote by Z^0 and $\tilde{Z}$ the matrices of general terms $Z^0_{i,j} := M^{\theta_0}_{\sigma_0(i)\sigma_0(j)}$ and $\tilde{Z}_{ij} := M^{\tilde{\theta}}_{\tilde{\sigma}(i)\tilde{\sigma}(j)}$ respectively, with $\tilde{\sigma}, \tilde{\theta}$ given by (17) and $i \neq j$. One interpretation of (17) if that the matrix $\tilde{Z}$ provides the best fit to the data X_{ij} with respect to the squared L^2 loss, when optimising over $\Sigma_e \times \Theta_n$.

As a first step, we show that $\tilde{Z}$ and Z^0 are close with high probability, a result in the spirit of Gao et al [16], Theorem 2.1. This follows from Lemma 2 below, which states that $\|\tilde{Z} - Z^0\|^2 \leq Cn \log k$ with probability at least $1 - e^{-n \log k}$, where $\|\cdot\|$ is the Frobenius norm.

In a second step, denoting $S_I^0 := \sigma_0^{-1}(\{1,2\})$ and recalling from (18) that $\tilde{S}_I = \tilde{\sigma}^{-1}(\{1,2\})$, we show that $\tilde{S}_I$ is close to S_I^0. To do so, one separately bounds from below some terms from the quantity $\|Z^0 - \tilde{Z}\|^2 = \sum_{i,j}(Z^0_{i,j} - \tilde{Z}_{i,j})^2$, recalling that one extends the Z matrices by symmetry and sets the diagonal to 0. First, using the definitions, (21), $|\theta| \leq \kappa$, and $n \geq 2$,

$$\sum_{i,j \in S_I^0 \setminus \tilde{S}_I} (Z^0_{ij} - \tilde{Z}_{ij})^2 \geq \frac{\kappa^2}{2} |S_I^0 \setminus \tilde{S}_I|^2,$$

if $|S_I^0 \setminus \tilde{S}_I| \geq 2$ (otherwise the inequality below holds trivially), as well as

$$\sum_{i \in S_I^0 \cap \tilde{S}_I\,, j \in S_I^0 \setminus \tilde{S}_I} (Z^0_{ij} - \tilde{Z}_{ij})^2 \geq \frac{\kappa^2}{2} |S_I^0 \cap \tilde{S}_I|\, |S_I^0 \setminus \tilde{S}_I|,$$

and, with $a \wedge b = \min(a,b)$, if $|S_I^0 \cap \tilde{S}_I| \geq 2$,

$$\sum_{i,j \in S_I^0 \cap \tilde{S}_I} (Z^0_{ij} - \tilde{Z}_{ij})^2 \geq \frac{1}{2} \left[(\theta_0 - \tilde{\theta})^2 \wedge (\theta_0 + \tilde{\theta})^2 \right] |S_I^0 \cap \tilde{S}_I|^2.$$

The previous bound on $\|\tilde{Z} - Z^0\|^2$ implies that

$$\max\left(|S_I^0 \setminus \tilde{S}_I|^2\,,\, |S_I^0 \cap \tilde{S}_I|\, |S_I^0 \setminus \tilde{S}_I| \right) \leq 2\kappa^{-2} n \log k \qquad \text{with high probability.}$$

It now follows from (22) that for any $\delta > 0$, one has $2\kappa^{-2} n \log k \leq \delta n^2/k^2$, provided d is small enough. So for small d, as $|S_I^0| = |S_I^0 \cap \tilde{S}_I| + |S_I^0 \setminus \tilde{S}_I|$, and as by assumption $\sigma_0 \in \Sigma_e$ so that $|S_I^0| \asymp n/k$, one deduces $|S_I^0 \cap \tilde{S}_I| \gtrsim n/k$. From the bound on $\|\tilde{Z} - Z^0\|^2$, it follows that

$$(\theta_0 - \tilde{\theta})^2 \wedge (\theta_0 + \tilde{\theta})^2 \lesssim \frac{n \log k}{n^2 k^{-2}} \lesssim \frac{k^2 \log k}{n} \qquad \text{with high probability.}$$

By (22) this shows that $\tilde{\theta}$ is close to either θ_0 or $-\theta_0$ up to $k(\log n/n)^{1/2} =: \rho \leq \kappa/2$. Now for any i, j in $\tilde{S}_I \setminus S_I^0$, if $\tilde{Z}_{ij} = \frac{1}{2} + \tilde{\theta}$ and $\tilde{\theta}$ is close to θ_0 (the cases where $\tilde{Z}_{ij} = \frac{1}{2} - \tilde{\theta}$ or $\tilde{\theta}$ is close to $-\theta_0$ are treated similarly), setting $Z^0_{ij} = c_{ij}$ and using (22), with $a_0 = 1/2$,

$$\begin{aligned} |Z^0_{ij} - \tilde{Z}_{ij}| &= |c_{ij} - (a_0 + \tilde{\theta})| = |c_{ij} - (a_0 + \theta^0) - (\tilde{\theta} - \theta^0)| \\ &\geq |c_{ij} - (a_0 + \theta^0)| - |\tilde{\theta} - \theta_0| \geq \kappa - \rho \geq \kappa/2. \end{aligned}$$

Therefore,

$$\kappa^2 |\tilde{S}_I \setminus S_I^0|^2 \lesssim \sum_{i,j \in \tilde{S}_I \setminus S_I^0} (Z^0_{ij} - \tilde{Z}_{ij})^2 \leq n \log k \qquad \text{with high probability}$$

and

$$\kappa^2 |S_I^0 \cap \tilde{S}_I|\, |\tilde{S}_I \setminus S_I^0| \lesssim \sum_{\substack{i \in S_I^0 \cap \tilde{S}_I \\ j \in \tilde{S}_I \setminus S_I^0}} (Z^0_{ij} - \tilde{Z}_{ij})^2 \leq n \log k \qquad \text{with high probability.}$$

Combining the previous bounds on cardinalities and denoting $A \,\Delta\, B := (A \setminus B) \cup (B \setminus A)$ for two sets A and B, one obtains

$$|\tilde{S}_I \,\Delta\, S_I^0|^2 + |\tilde{S}_I \,\Delta\, S_I^0||S_I^0 \cap \tilde{S}_I| \lesssim \kappa^{-2} n \log k \qquad \text{with high probability,}$$

which in turn implies that

$$|\{\tilde{S}_I \times \tilde{S}_I\} \,\Delta\, \{S_I^0 \times S_I^0\}| \lesssim \kappa^{-2} n \log k \qquad \text{with high probability.} \tag{31}$$

In a third and last step, we follow the proof of Theorem 2. Let $\hat{\sigma} = \hat{\sigma}_I$ be the mapping in (19). It is a map $\tilde{S}_I \to \{1,2\}$. Let $\bar{\sigma}$ be the mapping $S_I^0 \to \{1,2\}$ that coincides with $\hat{\sigma}$ on $\tilde{S}_I \cap S_I^0$ and with σ_0 on $S_I^0 \setminus \tilde{S}_I$. By definition we have, with $\Delta_n := \kappa^{-2} n \log k$,

$$\begin{aligned} Z_n(\hat{\sigma}, \tilde{S}_I, X) &= \sum_{i<j, i, j \in \tilde{S}_I} (-1)^{\mathbb{1}_{\hat{\sigma}(i)=\hat{\sigma}(j)}} \left(\frac{1}{2} - X_{ij}\right) \\ &= \sum_{i<j, i, j \in S_I^0} (-1)^{\mathbb{1}_{\bar{\sigma}(i)=\bar{\sigma}(j)}} \left(\frac{1}{2} - X_{ij}\right) + O(\Delta_n) \\ &= \delta(\bar{\sigma}, \sigma_0)\theta_0 + R_n(\bar{\sigma}) + O(\Delta_n), \end{aligned}$$

where for the second identity we have used that the X_{ij}s are bounded by 1 and (31), and $R_n(\sigma)$ is defined as in the proof of Theorem 2. Similarly, we have

$$Z_n(\sigma_0, \tilde{S}_I, X) = n_k \theta_0 + O(\Delta_n) + R_n(\sigma_0).$$

Let v_k and t_k be two sequences depending on n and k to be determined below. If $Z_n(\hat{\sigma}, \tilde{S}_I, X) \geq 0$, then $Z_n(\hat{\sigma}, \tilde{S}_I, X) \geq |Z_n(\sigma_0, \tilde{S}_I, X)| \geq Z_n(\sigma_0, \tilde{S}_I, X)$. Let Σ_0 be the set of all maps $S_0 \to \{1,2\}$. Then

$$\begin{aligned} &P_{\theta_0}\left[v_k|\hat{\theta} - \theta_0| \geq t, Z_n(\hat{\sigma}, \tilde{S}_I, X) \geq 0\right] \mathbb{1}_{\theta_0 > t_k} \\ &\leq P_{\theta_0}\left[\frac{v_k}{n_k}|Z_n(\hat{\sigma}, \tilde{S}_I, X) - Z_n(\sigma_0, \tilde{S}_I, X) - O(\Delta_n) - R_n(\sigma_0)| \geq t,\ Z_n(\hat{\sigma}, \tilde{S}_I, X) \geq 0\right] \mathbb{1}_{\theta_0 > t_k} \\ &\leq P_{\theta_0}\left[\frac{v_k}{n_k}\left[Z_n(\hat{\sigma}, \tilde{S}_I, X) - Z_n(\sigma_0, \tilde{S}_I, X)\right] + \frac{v_k}{n_k}\left[|R_n(\sigma_0)| + O(\Delta_n)\right] \geq t\right] \mathbb{1}_{\theta_0 > t_k} \\ &\leq P_{\theta_0}\left[\frac{v_k}{n_k}\theta_0(\delta(\bar{\sigma}, \sigma_0) - n_k) + \frac{v_k}{n_k}\left[R_n(\hat{\sigma}) + |R_n(\sigma_0)| + O(\Delta_n)\right] \geq t\right] \mathbb{1}_{\theta_0 > t_k} \\ &\leq P_{\theta_0}\left[2\frac{v_k}{n_k}\sup_{\sigma \in \Sigma_0}|R_n(\sigma)| + \frac{v_k}{n_k}O(\Delta_n) \geq t\right], \end{aligned}$$

where the last inequality uses $\delta(\bar{\sigma}, \sigma_0) - n_k \leq 0$, see Lemma 1 with n_k in place of b_n, and $\theta_0 \geq 0$.

Second, as $Z_n(\hat{\sigma}, \tilde{S}_I, X) < 0$ implies $Z_n(\hat{\sigma}, \tilde{S}_I, X) < -|Z_n(\sigma_0, \tilde{S}_I, X)|$ by definition of $\hat{\sigma}$,

$$\begin{aligned} &P_{\theta_0}\left[Z_n(\hat{\sigma}, \tilde{S}_I, X) < 0\right] \mathbb{1}_{\theta_0 > t_k} \\ &\leq P_{\theta_0}\left[Z_n(\hat{\sigma}, \tilde{S}_I, X) < -|Z_n(\sigma_0, \tilde{S}_I, X)|\right] \mathbb{1}_{\theta_0 > t_k} \\ &\leq P_{\theta_0}\left[Z_n(\hat{\sigma}, \tilde{S}_I, X) < -|\theta_0| n_k + O(\Delta_n) + |R_n(\sigma_0)|\right] \mathbb{1}_{\theta_0 > t_k} \\ &\leq P_{\theta_0}\left[\theta_0(\delta(\bar{\sigma}, \sigma_0) + n_k) - 2O(\Delta_n) < 2\sup_{\sigma \in \Sigma}|R_n(\sigma)|\right] \mathbb{1}_{\theta_0 > t_k} \\ &\leq P_{\theta_0}\left[2\sup_{\sigma \in \Sigma_0}|R_n(\sigma)| > \frac{t_k n_k}{8} - O(\Delta_n)\right], \end{aligned}$$

where for the last inequality we have used the first inequality of Lemma 1. Also,

$$\begin{aligned} &P_{\theta_0}\left[v_k|\hat{\theta} - \theta_0| \geq t\right] \mathbb{1}_{|\theta_0| \leq t_k} \\ &\leq \quad P_{\theta_0}\left[\frac{v_k}{n_k}|\theta_0||\delta(\bar{\sigma}, \sigma_0) - n_k| + \frac{v_k}{n_k}\left[O(\Delta_n) + |R_n(\bar{\sigma})| + |R_n(\sigma_0)|\right] \geq t\right] \mathbb{1}_{|\theta_0| \leq t_k} \\ &\leq \quad P_{\theta_0}\left[2\frac{v_k}{n_k}\sup_{\sigma \in \Sigma_0}|R_n(\sigma)| + \frac{v_k}{n_k}O(\Delta_n) \geq t - v_k t_k\right]. \end{aligned}$$

By the same argument as in the proof of Theorem 2, the supremum $\sup_{\sigma\in\Sigma_0}|R_n(\sigma)|$ is of the order $|\Sigma_0|^{3/2} \asymp (n/k)^{3/2}$, by definition of Σ_e. Recall that $n_k \asymp (n/k)^2$. Set $v_k = \sqrt{n/k}$ and $t_k = Dv_k^{-1}$, with D a large enough constant. Assumption (22) ensures that $\Delta_n = \kappa^{-2} n \log k = O((n/k)^{3/2})$. Hence by taking t a large enough constant, one obtains that the last three displays are bounded above by $e^{-Cn/k}$, which concludes the proof in the fixed design case, proceeding as in the proof of Theorem 2 to get the final bound in expectation.

The proof in the random design case is obtained by first deriving the results conditionally on φ and then integrating out φ, as we did in the proof of Theorem 2. The first part is almost identical to the fixed design case: one only needs to note that one can restrict to mappings φ that belong to, essentially, Σ_e. Denote by Σ_e' the subset of those $\sigma \in \Sigma_e$ satisfying $|\sigma^{-1}(j) - \frac{n}{k}| < \frac{n}{2k}$. Then

$$\begin{aligned} E_{\theta_0}[(\hat\theta - \theta_0)^2] &= E_{\theta_0}[(\hat\theta - \theta_0)^2 \mathbb{1}\{\varphi \in \Sigma_e'\}] + E_{\theta_0}[(\hat\theta - \theta_0)^2 \mathbb{1}\{\varphi \in \Sigma_e'\}] \\ &= E_{\theta_0}\left[E_{\theta_0}[(\hat\theta - \theta_0)^2 \,|\, \varphi]\, \mathbb{1}\{\varphi \in \Sigma_e'\}\right] + P_{\theta_0}[\varphi \notin \Sigma_e']. \end{aligned}$$

For the first term, one can apply the arguments above in fixed design, while for the second an application of Bernstein's inequality gives

$$P_{\theta_0}\left[\varphi \notin \Sigma_e'\right] \le kP\left[|\mathrm{Bin}(n,k) - \frac{n}{k}| > \frac{n}{2k}\right] \le 2ke^{-\frac{n}{10k}}.$$

By (22), $k^3 \le Cn$ holds for C large enough—note that κ must be smaller than 1, as the entries of M are in $[0,1]$. One deduces $ke^{-\frac{n}{10k}} \le \frac{k}{n} n e^{-dn^{2/3}} \le C\frac{k}{n}$, for d small enough and C large enough, so the quadratic risk is at most Ck/n in this case as well. □

Lemma 2. *Let $Z^0_{i,j} := M^{\theta_0}_{\sigma_0(i)\sigma_0(j)}$ and $\tilde Z_{ij} := M^{\tilde\theta}_{\tilde\sigma(i)\tilde\sigma(j)}$, with $\tilde\sigma, \tilde\theta$ given by (17). Let $\|\cdot\|$ denote the matrix Frobenius norm. With probability at least $1 - e^{-cn\log k}$,*

$$\|\tilde Z - Z^0\|^2 \lesssim Cn\log k.$$

Proof of Lemma 2. Let θ_1 denote the element of Θ_n closest to θ_0, so that $|\theta_0 - \theta_1| \le n^{-2}$. Let Z^1 be the matrix given by $Z^1_{ij} := M^{\theta_1}_{\sigma_0(i)\sigma_0(j)}$. By definition, $\|X - \tilde Z\|^2 \le \|X - Z^1\|^2$ and hence $\|Z^1 - \tilde Z\|^2 + 2\langle X - Z^1, Z^1 - \tilde Z\rangle \le 0$, so

$$\|Z^1 - \tilde Z\|^2 \le 2\|Z^1 - \tilde Z\| \sup_{\theta\in\Theta_n,\ \sigma\in\Sigma_e} |\langle X - Z^1, \frac{M^\theta_{\sigma(i)\sigma(j)} - Z^1}{\|M^\theta_{\sigma(i)\sigma(j)} - Z^1\|}\rangle|.$$

Denote $T_n(\sigma,\theta) := \langle X - Z^0, (M^\theta_{\sigma(i)\sigma(j)} - Z^0)/\|M^\theta_{\sigma(i)\sigma(j)} - Z^0\|\rangle$. By Hoeffding's inequality,

$$P[|T_n(\sigma,\theta)| > t] \le 2\exp\{-2t^2\}.$$

The cardinality of the set $\Theta_n \times \Sigma_e$ is bounded above by $(2n^2+1)k^n \lesssim k^{Cn}$. A union bound then shows that, with probability at least $1 - e^{-cn\log k}$,

$$\sup_{\theta\in\Theta,\ \sigma\in\Sigma_e} |T_n(\sigma,\theta)| \le C\sqrt{n\log k}.$$

Inserting this back into the previous inequality on $\|Z^1 - \tilde Z\|$ leads to $\|Z^1 - \tilde Z\| \le C\sqrt{n\log k}$ with probability at least $1 - e^{-cn\log k}$. As $\|Z^0 - Z^1\|^2 \le Cn^2/n^2 \le C$, the triangle inequality leads to the result. □

Appendix C: Proof of the lower bound for graphon functionals

We start with a preliminary result in the graphon model (2), assuming that the function f is piecewise constant, with different values taken along blocks corresponding to a regular partition of $[0,1]^2$ in $k \times k = k^2$ blocks, $k \geq 2$. That defines a law of the form

$$\frac{1}{k^n} \sum_{\varphi \in [k]^n} \bigotimes_{i<j} \mathrm{Be}(Q_{\varphi(i)\varphi(j)}) = P_{e_k,Q},$$

where φ is an element of $[k]^n$ and $Q = Q^\theta$ a given $k \times k$ matrix defined below. In the next statement and proof, E_θ denotes the expectation under this distribution. Denote by O_k the $k \times k$ matrix with only ones as coefficients,

$$O_k = \begin{bmatrix} 1 & \cdots & 1 \\ \vdots & \vdots & \vdots \\ 1 & \cdots & 1 \end{bmatrix},$$

and, for a *symmetric* $l \times l$ matrix A with coefficients $A_{ij} \in [0,1]$, define the $k \times k$ matrix

$$B := \left[\begin{array}{c|c} A & -A \\ \hline -A & A \end{array}\right].$$

We define $Q = Q^\theta$ as the $k \times k = (2l) \times (2l)$ matrix

$$Q = Q^\theta = \frac{1}{2} \cdot O_k + \theta \cdot B = \frac{1}{2} \cdot O_k + \theta \left[\begin{array}{c|c} A & -A \\ \hline -A & A \end{array}\right]. \tag{32}$$

Lemma 3. *Let $k = 2l$ be an even integer and A an arbitrary symmetric $l \times l$ matrix. Let $Q = Q^\theta$ be the matrix defined in (32). There exists a constant $c_3 > 0$ such that*

$$\inf_{\hat\theta} \sup_{\theta \in (-1/2, 1/2)} E_\theta(\hat\theta(X) - \theta)^2 \geq \frac{c_3}{n},$$

where the infimum is over all estimators of θ valid under $E_\theta = E_{P_{e_k,Q^\theta}}$.

Proof of Lemma 3. Let $\mathbb{Q} = P_{e_k,Q^\theta}$ be as above. That is,

$$\mathbb{Q} = \frac{1}{k^n} \sum_{\varphi \in [k]^n} P_{Z^\theta_\varphi}, \qquad P_{Z^\theta_\varphi} = \bigotimes_{i<j} \mathrm{Be}(Q^\theta_{\varphi(i)\varphi(j)}),$$

with $\{Z^\theta_\varphi\}$ the matrix of general term $z_{ij}(\varphi,\theta) = Q^\theta_{\varphi(i)\varphi(j)}$, for φ ranging over the set $[k]^n$. Let $\mathbb{P}$ denote the Erdös-Renyi $ER(1/2)$ distribution over n nodes, which also corresponds to P_{e_k,Q^θ} for $\theta = 0$. Consider the functional ψ defined as,

$$\psi(P_{e_k,Q^\theta}) = \theta.$$

By definition, for any $\varphi \in [k]^n$, we have $\psi(\mathbb{P}) = \psi(P_{e_k,Q^0}) = 0$ and $\psi(\mathbb{Q}) = \psi(P_{e_k,Q^\theta}) = \theta$. The same computation as in the proof of Theorem 1 now shows that, for B given in the display below,

$$\|\mathbb{P} - \mathbb{Q}\|_1^2 \leq \frac{1}{k^{2n}} \sum_{\varphi,\psi \in [k]^n} \exp\Big\{4\theta^2 \sum_{i<j} B_{\varphi(i)\varphi(j)} B_{\psi(i)\psi(j)}\Big\} - 1.$$

The last term in the bound can be interpreted as an expectation over φ, ψ, where both variables are sampled uniformly from the set of all mappings from $\{1,\ldots,n\}$ to $\{1,\ldots,k\}$. Recall that $l = k/2$ and for any integer s, denote by $[s]_l$ the integer in $\{1,\ldots,l\}$ that equals s modulo l, plus 1. The variable $B_{\varphi(i)\varphi(j)}$ can be written

$$B_{\varphi(i)\varphi(j)} = (-1)^{\varphi(i)>l}(-1)^{\varphi(j)>l} A_{[\varphi(i)]_l[\varphi(j)]_l}.$$

When φ follows the uniform distribution over $[k]^n$, the variables $(\varphi(i))_i$ are independent and are marginally uniform over $[k]$. Also, the variables $((-1)^{\varphi(i)>l})_i$ and $([\varphi(i)]_l)_i$ are independent under the uniform distribution for φ as we show next. If Pr denotes the corresponding distribution, then for any $i \leq n$ and any $s \leq l$

$$\Pr\Big[\varphi(i) > l, [\varphi(i)]_l = s\Big] = \frac{1}{k} = \frac{1}{2} \cdot \frac{1}{l} = \Pr[\varphi(i) > l]\Pr[\varphi(i)]_l = s].$$

Note the identity holds both for $k \leq n$ and $k > n$. Set $R_i := (-1)^{\varphi(i)>l}$ and $a_{ij} := A_{[\varphi(i)]_l[\varphi(j)]_l}$. Deduce from the previous reasoning that the variables $(R_i)_i$ and $(a_{ij})_{i<j}$ are independent. Now, denoting by E the expectation under Pr,

$$\begin{aligned}\|\mathbb{P} - \mathbb{Q}\|_1^2 &\leq E\left[\exp\left\{4\theta^2 \sum_{i<j} B_{\varphi(i)\varphi(j)} B_{\psi(i)\psi(j)}\right\}\right] - 1 \\ &\leq E\left[E\left[\exp\left\{4\theta^2 \sum_{i<j} a_{ij} R_i R_j\right\}\middle| a_{ij}\right]\right] - 1.\end{aligned}$$

As $(R_i)_i$ and $(a_{ij})_{i<j}$ are independent, one can compute the inner expectation in the last display under the distribution of $(R_i)_i$, the a_{ij}'s being fixed. The $(R_i)_i$ form a sample of independent Rademacher variables, hence

$$Z_n := \sum_{i<j} a_{ij} R_i R_j$$

is a Rademacher chaos of order 2 with weights (a_{ij}). Suppose the matrix A is not identically zero (otherwise the bound below holds trivially). By Lemma 4, for any $c > 1$ one can find $\lambda > 0$ with

$$E\exp\Big[\frac{|Z_n|}{\lambda\|Z_n\|_2}\,\Big|\, a_{ij}\Big] \leq c \qquad \text{where} \qquad \|Z_n\|_2^2 = E[|Z_n|^2 \,|\, a_{ij}] = \sum_{i<j} a_{ij}^2.$$

Choose $c = 3/2$. By definition, all a_{ij}s are bounded by 1. There is hence a $\lambda > 0$ such that, if $\theta^2 = 2/(\lambda n)$,

$$E\exp\big[4\theta^2|Z_n|\ \big|\, a_{ij}\big] - 1 \leq 3/2 - 1 = 1/2.$$

The result now follows from an application of Lemma 6 to the functional ψ. □

Proof of Theorem 5. Let us recall the definition, for any $0 \leq \theta \leq 1$, of the function $w = w_\theta$ in (25)

$$w_\theta(x, y) = \frac{1}{2} - \theta(x - \frac{1}{2})(y - \frac{1}{2})$$

and let $\langle w_\theta \rangle$ be its graphon equivalence class. By definition, $\langle w_\theta \rangle$ belongs to $\mathcal{P}$. One has

$$\int_{[0,1]^2} w_\theta(x, y)^2 dxdy = \frac{1}{4} + c^2\theta^2,$$

for some constant $c > 0$, so that $\tau(\langle w_\theta \rangle) - \tau(\langle w_0 \rangle) = c\theta$. The function w_0 is the constant $1/2$, and the density of the data distribution $P_{\langle w \rangle}$ with respect to counting measure on $\{0, 1\}^{n(n-1)/2}$ is

$$p_w(\{x_{ij}\}_{i<j}) = \int \cdots \int \prod_{i<j} \mathrm{Be}(w(u_i, u_j))(x_{ij}) du_1 \cdots du_n,$$

where, for any u, v, z in $[0, 1]$, we have set

$$\mathrm{Be}(z)(x_{ij}) = z^{x_{ij}}(1 - z)^{1-x_{ij}}.$$

Next one shows that $P_{\langle w\rangle}$ is close in the total variation sense to a discrete mixture of the previous Bernoulli-probability distributions, provided the number of points in the mixture is suitably large. To do so, we approximate the function P_n defined by

$$P_n : (u_1, \dots, u_n) \mapsto \prod_{i<j} \mathrm{Be}(w(u_i, u_j))(x_{ij}),$$

by a piecewise constant function $h_{N,\theta} = h_N$, where $[0,1]^n$ is split into N^n blocks, $N \geq 1$, using a regular grid of $[0,1]^n$ with points $(i_1/N, \dots, i_n/N)$ and $0 \leq i_j \leq N$ for all j. To do so, one just replaces $w(u_i, u_j)$ by, say, the value of w on the middle of the block the point (u_i, u_j) belongs to. This defines a function

$$Q_{n,N} : (u_1, \dots, u_n) \mapsto \prod_{i<j} \mathrm{Be}(\bar{w}(u_i, u_j))(x_{ij}),$$

where $\bar{w}$ is constant on every block of the subdivision. Let Q_w^N denote the corresponding measure, with density

$$q_w^N(\{x_{ij}\}_{i<j}) = \int \cdots \int \prod_{i<j} \mathrm{Be}(\bar{w}(u_i, u_j))(x_{ij}) du_1 \cdots du_n.$$

Taking $w = w_\theta$ as above, the function P_n is a polynomial in $u_1, \dots, u_n$, and its degree with respect to each variable u_i is $n-1$. The partial derivatives of P_n can be computed, and each of them can be seen to be bounded by $n-1$: For each variable, only $n-1$ non-zero terms appear when evaluating the partial derivative, and each term is uniformly bounded by 1. Consequently, if $(u_1, \dots, u_n)$ and $(u'_1, \dots, u'_n)$ belong to the same block,

$$|P_n(u_1, \dots, u_n) - P_n(u'_1, \dots, u'_n)| \leq (n-1) \sum_{i=1}^{n} |u_i - u'_i| \leq n^2/N.$$

For $w = w_\theta$ as above, we can thus bound the total variation distance as

$$\begin{aligned} \|P_{\langle w\rangle} - Q_w^N\|_1 &\leq \sum_{x \in \{0,1\}^{\frac{n(n-1)}{2}}} |p_w - q_w^N|(x) \\ &\leq n^2 \max_{x \in \{0,1\}^{\frac{n(n-1)}{2}}} |p_w - q_w^N|(x) \leq n^2(n^2/N) = n^4/N. \end{aligned}$$

Each probability measure Q_w^N is a mixture of N^n distributions, each of which in turn corresponds to a block in the subdivision of $[0,1]^n$. One can rewrite

$$Q_w^N = \frac{1}{N^n} \sum_{\varphi \in [N]^n} \bigotimes_{i<j} \mathrm{Be}(M_{\varphi(i)\varphi(j)}),$$

where the matrix $M = (M_{pq})_{1 \leq p,q \leq N}$ is the symmetric matrix with terms

$$M_{pq} = w_\theta\left(\frac{p - \frac{1}{2}}{N}, \frac{q - \frac{1}{2}}{N}\right).$$

If N is even, which one can assume without loss of generality, the matrix M is exactly of the same form as Q in (32), with elements in $(0,1)$, so one can use the bound in $\|\cdot\|_1$-distance between measures obtained in the proof of Lemma 3. Note that the argument remains valid even if the number of classes exceeds the number of observations n, which will be of importance below. For a small constant c and $\theta^2 = \kappa/n$, we obtain

$$\|Q_{w_\theta}^N - Q_{w_0}^N\|_1 \leq c$$

for κ sufficiently small. Choosing $N = Cn^4$, for $C > 0$ large enough, leads to

$$\begin{aligned}\|P_{\langle w_\theta\rangle} - P_{\langle w_0\rangle}\|_1 &\leq \|P_{\langle w_\theta\rangle} - Q^N_{w_\theta}\|_1 + \|Q^N_{w_\theta} - Q^N_{w_0}\|_1 + \|Q^N_{w_0} - P_{\langle w_0\rangle}\|_1 \\ &\leq n^4/N + c + 0 \leq c' < 1/2.\end{aligned}$$

An application of Lemma 6 with the functional $\psi(P_{\langle w\rangle}) := \psi(\langle w\rangle)$ concludes the proof of the lower bound in Theorem 5 in the case where $\psi(\cdot) = \tau(\cdot)$. The lower bound for a general ψ follows by the same proof, noting that the specific form of the functional only comes in through the difference $\psi(\langle w_\theta\rangle) - \psi(\langle w_0\rangle)$, which behaves as for $\tau(\cdot)$ by assumption.

For the upper-bound, we first link the squared distance to the truth for the functional to the squared L^2-distance of corresponding graphons. Let w, w_1 be two fixed graphon functions, and suppose that at least one of these is non constant (almost everywhere), which means that either $\tau(\langle w\rangle) > 0$ or $\tau(\langle w_1\rangle) > 0$. Then, writing simply $\int$ to denote the double integral on $[0,1]^2$,

$$\begin{aligned}\tau(\langle w_1\rangle) - \tau(\langle w\rangle) &= \left\{\int (w_1 - \int w_1)^2\right\}^{1/2} - \left\{\int (w - \int w)^2\right\}^{1/2} \\ &= \frac{\int (w_1 - \int w_1)^2 - \int (w - \int w)^2}{\left\{\int (w_1 - \int w_1)^2\right\}^{1/2} + \left\{\int (w - \int w)^2\right\}^{1/2}},\end{aligned}$$

where the denominator is nonzero by assumption on w, w_1; we henceforth denote it c. Then

$$\begin{aligned}\tau(\langle w_1\rangle) - \tau(\langle w\rangle) &\leq c^{-1} \int \left\{w_1 - w - \int (w_1 - w)\right\}\left\{w_1 + w - \int (w_1 + w)\right\} \\ &\leq c^{-1} \left[\int \left\{w_1 - w - \int (w_1 - w)\right\}^2\right]^{1/2} \left[\int \left\{w_1 + w - \int (w_1 + w)\right\}^2\right]^{1/2}.\end{aligned}$$

The two factors in brackets are bounded as follows: For the second term, apply the inequality $(a+b)^2 \leq 2a^2 + 2b^2$, followed by $\sqrt{u+v} \leq \sqrt{u} + \sqrt{v}$, which yields

$$\left[\int \left\{w_1 + w - \int (w_1 + w)\right\}^2\right]^{1/2} \leq \sqrt{2}\left[\int (w_1 - \int w_1)^2 + \int (w - \int w)^2\right]^{1/2} \leq \sqrt{2}c.$$

For the first term, use $(a+b)^2 \leq 2a^2 + 2b^2$ and $(\int w)^2 \leq \int w^2$. That yields

$$\{\tau(\langle w_1\rangle) - \tau(\langle w\rangle)\}^2 \leq 8 \int (w_1 - w)^2,$$

and this inequality clearly still holds true in case $\tau(\langle w\rangle) = \tau(\langle w_1\rangle) = 0$. One concludes that

$$\{\tau(\langle w_1\rangle) - \tau(\langle w\rangle)\}^2 \leq 8 \inf_{T\in\mathcal{T}} \int\int_{[0,1]^2} |w_1(T(x), T(y)) - w(x,y)|^2 \, dxdy =: \delta^2(w_1, w),$$

where $\mathcal{T}$ is the set of all measure-preserving bijections of $[0,1]$. Indeed, the previous inequalities hold true for any choice of representer of the graphon w_1, so one can take the infimum over T in the previous bounds. By Proposition 3.5 of [23], for data X generated from P_w, there exists an estimator $\hat{w} = \hat{w}(X)$ that satisfies $E_{P_w}[\delta^2(\hat{w}, w)] \leq Cn^{-1}$. Since $\langle w\rangle$ has a representer that belongs to $\mathcal{P}_B$ by assumption, it belongs in particular to the Hölder class $\Sigma(1, L)$, provided L is chosen large enough. For the plug-in estimator $\hat{\tau}(X) := \tau(\hat{w})$, combining the previous result with the last display implies

$$E\left[\{\hat{\tau}(X) - \tau(\langle w\rangle)\}^2\right] \leq \frac{C}{n},$$

for C large enough depending only on $\mathcal{P}_B$, which concludes the proof. □

Appendix D: Useful lemmas

Let $\{Z_i,\ i \geq 1\}$ be i.i.d. Rademacher variables. For reals x_{ij} and $N \geq 2$, set

$$Y := Y_N = \sum_{i<j\leq N} y_{ij} Z_i Z_j,$$
$$s(Y)^2 = \sum_{i<j\leq N} y_{ij}^2.$$

Lemma 4 (Corollary 3.2.6 in de la Peña and Giné [13])**.** *Let $N \geq 2$, and $Y = Y_N$ and $s(Y)$ as above. For every $c > 1$, there exists $\lambda = \lambda(c) > 0$ independent of N such that*

$$\mathbb{E} \exp\left[\frac{|Y|}{\lambda s(Y)}\right] \leq c.$$

We repeatedly use Lemma 4 in the case where all y_{ij} are equal to 1, for various values of N. In such a setting, a reformulation is as follows. For any $c > 1$ and $N \geq 2$, one can find a constant $a = a(c)$ independent of N such that

$$\mathbb{E} \exp\left[a\frac{|Y_N|}{N}\right] \leq c. \tag{33}$$

Lemma 5 (Rademacher chaos with explicit constant)**.** *Let $N \geq 2$, and $Y = Y_N$ and $s(Y)$ as above. For any $0 \leq \delta \leq 1$,*

$$\mathbb{E} \exp\left[\delta\frac{|Y|}{s(Y)}\right] - 1 \leq r(\delta), \quad \textit{with} \quad r(\delta) = \delta + \frac{\delta^2}{2} + \frac{8\delta^3}{6} + \frac{1}{1-e\delta}\frac{(e\delta)^4}{\sqrt{8\pi}}.$$

The lemma applied with $\delta = 1/3$ gives a bound 1.87 for the right hand side.

Proof. Theorem 3.2.2 in de la Peña and Giné [13] gives, for any $k \geq 2$,

$$E|Y|^k \leq (k-1)^k s(Y)^k.$$

For $k = 1$ one has $E[|Y|] \leq E[Y^2]^{1/2} = s(Y)$. From this one deduces that for any $0 \leq \delta \leq 1$,

$$E \exp\left[\delta\frac{|Y|}{s(Y)}\right] \leq 1 + \delta + \frac{\delta^2}{2} + \frac{8\delta^3}{6} + \sum_{k\geq 4}\frac{(k-1)^k}{k!}\delta^k,$$

and the result follows from an application of the nonasymptotic Stirling bound $k! \geq e^{-k}k^{k+\frac{1}{2}}\sqrt{2\pi}$ valid for $k \geq 1$. □

Lemma 6 (Le Cam's method 'point versus mixture')**.** *Let $\mathcal{P} = \{P_M,\ M \in \mathcal{M}\}$ be a collection of probability measures indexed by an arbitrary set $\mathcal{M} = \{M_0, M_1, \ldots, M_N\}$, $N > 1$. Set*

$$\mathbb{P} = P_{M_0}, \quad \mathbb{Q} = \frac{1}{N}\sum_{k=1}^{N} P_{M_k}.$$

If ψ is a real-valued functional such that $\psi(P_{M_0}) = \tau$ and $\psi(P_{M_i}) = \theta$ for any $i = 1, \ldots, N$, then

$$\inf_{\hat{\theta}} \sup_{M\in\mathcal{M}} E_{P_M}(\hat{\psi}(X) - \psi(P_M))^2 \geq \frac{1}{4}(\theta-\tau)^2\left(1 - \frac{1}{2}\|\mathbb{P} - \mathbb{Q}\|_1\right),$$

where the infimum is over all estimators $\hat{\psi}(X)$ of $\psi(P_M)$ based on the observation of $X \sim P_M$.

Proof. This is a standard variation on the case where $N = 1$ stated in e.g. [31]. □

Lemma 7 (Bound on total variation distance)**.** *For $n \geq 1$, let $P_1, \ldots, P_n$ and $Q_1(k), \ldots, Q_n(k)$ for $1 \leq k \leq N$, for some $N \geq 1$, be probability measures. Set*

$$P = \bigotimes_{i=1}^{n} P_i, \quad Q(k) = \bigotimes_{i=1}^{n} Q_i(k) \quad \mathbb{P} = P, \quad \mathbb{Q} = \frac{1}{N}\sum_{k=1}^{N} Q(k).$$

Suppose that for any i, $Q_i(k)$ has density $1+\Delta_i(k)$ with respect to P_i. Denote $\tau_i(k,l) = P_i\Delta_i(k)\Delta_i(l)$. Then, for $\chi^2(\mathbb{Q},\mathbb{P}) = \int (d\mathbb{Q}/d\mathbb{P} - 1)^2 d\mathbb{P}$,

$$\|\mathbb{P} - \mathbb{Q}\|_1^2 \leq \chi^2(\mathbb{Q},\mathbb{P}) = \frac{1}{N^2}\sum_{k,l}\prod_{i=1}^{n}\{1 + \tau_i(k,l)\} - 1.$$

Proof. The first bound on distances is standard, while the second bound follows from elementary calculations. □

Lemma 8. *Let N, R be two integers with $N \geq 2$, $1 \leq R \leq N$, and $(P_i)_{i\in I}$ be an arbitrary collection of probability measures with $|I| = N$. If $J \subset I$ and $|J| = R$, we have*

$$\left\| \frac{1}{N}\sum_{i\in I} P_i - \frac{1}{R}\sum_{i\in J} P_i \right\|_1 \leq 2\left(1 - \frac{R}{N}\right).$$

Proof. The result follows by spliting the sum over I in a sum over J and $I \setminus J$, applying the triangle inequality and using the fact that $\| \sum_{j\in J} P_j \|_1 = |J|$. □

Lemma 9. *Let $(\lambda_i)_{1\leq i\leq N}$ be given weights in $[0,1]$ such that $\sum_i \lambda_i = 1$. Let $M_1 := \sum_{i=1}^N \lambda_i P_i$ and $M_2 := \sum_{i=1}^N \lambda_i Q_i$ be two mixtures and R a given probability measure on the same space. We have*

$$KL(M_1, M_2) \leq \sum_{i=1}^{N} \lambda_i KL(P_i, Q_i)$$

as well as

$$KL(R, M_2) \leq \sum_{i=1}^{N} \lambda_i KL(R, Q_i).$$

Proof. Let μ be a common dominating measure for all probability measures R, P_i, Q_i, $1 \leq i \leq N$, and let r, p_i, q_i the corresponding densities. By definition

$$KL(M_1, M_2) = \int \log\left[\frac{\sum_i \lambda_i p_i}{\sum_i \lambda_i q_i}\right]\left[\sum_i \lambda_i p_i\right] d\mu$$

and

$$KL(R, M_2) = \int \log\left[\frac{r}{\sum_i \lambda_i q_i}\right] r d\mu.$$

The result follows by convexity of the mappings $(x,y) \to x\log(x/y)$ and $y \to x\log(x/y)$ on $(0,\infty)^2$ and $(0,\infty)$ respectively. □

References

[1] E. S. Allman, C. Matias, and J. A. Rhodes. Parameter identifiability in a class of random graph mixture models. *J. Statist. Plann. Inference*, 141(5):1719–1736, 2011.

[2] C. Ambroise and C. Matias. New consistent and asymptotically normal parameter estimates for random-graph mixture models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 74(1):3–35, 2012.

[3] E. Arias-Castro, E. J. Candès, and A. Durand. Detection of an anomalous cluster in a network. *Ann. Statist.*, 39(1):278–304, 2011.
[4] E. Arias-Castro and N. Verzelen. Community detection in dense random networks. *Ann. Statist.*, 42(3):940–969, 2014.
[5] P. Bickel and A. Chen. A nonparametric view of network models and Newman-Girvan and other modularities. *PNAS*, 106(50):21068–21073, 2009.
[6] P. Bickel and H. Chernoff. Asymptotic distribution of the likelihood ratio statistic in a prototypical non regular problem. In *Statistics and Probability: A Raghu Raj Bahadur Festschrift*, pages 83–96. Wiley, New York, 1993.
[7] P. Bickel, D. Choi, X. Chang, and H. Zhang. Asymptotic normality of maximum likelihood and its variational approximation for stochastic blockmodels. *Ann. Statist.*, 41(4):1922–1943, 2013.
[8] P. J. Bickel, A. Chen, and E. Levina. The method of moments and degree distributions for network models. *Ann. Statist.*, 39(5):2280–2301, 2011.
[9] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós, and K. Vesztergombi. Convergent sequences of dense graphs. I. Subgraph frequencies, metric properties and testing. *Adv. Math.*, 219(6):1801–1851, 2008.
[10] C. Butucea and Y. I. Ingster. Detection of a sparse submatrix of a high-dimensional noisy matrix. *Bernoulli*, 19(5B):2652–2688, 2013.
[11] A. Celisse, J.-J. Daudin, and L. Pierre. Consistency of maximum-likelihood and variational estimators in the stochastic block model. *Electron. J. Stat.*, 6:1847–1899, 2012.
[12] A. Channarond, J.-J. Daudin, and S. Robin. Classification and estimation in the stochastic blockmodel based on the empirical degrees. *Electron. J. Stat.*, 6:2574–2601, 2012.
[13] V. H. de la Peña and E. Giné. *Decoupling*. Probability and its Applications. Springer-Verlag, New York, 1999.
[14] D. Easley and J. Kleinberg. *Networks, crowds and markets*. Cambridge University Press, 2010.
[15] S. Gadat, C. Marteau, and C. Maugis-Rabusseau. Parameter recovery in two-component contamination mixtures: the L2 strategy. 2016. preprint.
[16] C. Gao, Y. Lu, and H. H. Zhou. Rate-optimal graphon estimation. *Ann. Statist.*, 43(6):2624–2652, 2015.
[17] A. Goldenberg, A. X. Zheng, S. E. Fienberg, and E. M. Airoldi. A survey of statistical network models. *Foundations and Trends in Machine Learning*, 2(2):129–233, 2009.
[18] M. S. Handcock and K. J. Gile. Modeling social networks from sampled data. *Ann. App. Statist.*, 4:5–25, 2010.
[19] J. A. Hartigan. A failure of likelihood asymptotics for normal mixtures. In *Proceedings of the Berkeley conference in honor of Jerzy Neyman and Jack Kiefer, Vol. II (Berkeley, Calif., 1983)*, Wadsworth Statist./Probab. Ser., pages 807–810. Wadsworth, Belmont, CA, 1985.
[20] P. D. Hoff, A. E. Raftery, and M. S. Handcock. Latent space approaches to social network analysis. *J. Amer. Statist. Assoc.*, 97(460):1090–1098, 2002.
[21] S. Janson. *Graphons, cut norm and distance, couplings and rearrangements*. New York Journal of Mathematics. NYJM Monographs. 2013.
[22] J. Kahn and P. Heinrich. Optimal rates for finite mixture estimation. 2015. preprint.
[23] O. Klopp, A. Tsybakov, and N. Verzelen. Oracle inequalities for network models and sparse graphon estimation. 2015. preprint.
[24] E. D. Kolaczyk. *Statistical Analysis of Network Data*. Springer, 2009.
[25] E. Mossel, J. Neeman, and A. Sly. Consistency thresholds for the planted bisection model. *Electron. J. Probab.*, 21:1–24, 2016.
[26] M. Newman. *Networks. An Introduction*. Oxford University Press, 2009.
[27] P. Orbanz and D. M. Roy. Bayesian models of graphs, arrays and other exchangeable random structures. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 37:437–461, 2015.
[28] A. van der Vaart and S. van der Pas. Bayesian community detection. 2016. preprint.
[29] N. Verzelen and E. Arias-Castro. Community detection in sparse random networks. *Ann. Appl. Probab.*, 25(6):3465–3510, 2015.

[30] P. Wolfe and S. Ohlede. Nonparametric graphon estimation. 2013. preprint.
[31] B. Yu. Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer, New York, 1997.
[32] A. Zhang and H. Zhou. Minimax rates of community detection in stochastic block models. 2015. preprint.